INTERACTIVE BROCHURE

White paper

Rethinking power architecture in large-scale data centres

How block-redundant (Catcher) architectures optimise power infrastructure while maintaining high availability.

discover ›



When **energy** matters

socomec
Innovative Power Solutions

Contents

1. Introduction
2. Design principles
3. Distributed architecture
4. Block-redundant architecture (Catcher)
5. User case





The challenge is no longer only to ensure resilience, but to structure power architecture in such a way that maintains availability while controlling infrastructure growth, complexity and costs - at scale.

Power architecture in the era of large-scale data centres

High availability remains a non-negotiable requirement in data centres. Ensuring a continuous power supply to IT loads is fundamental, as any interruption directly impacts service continuity and operational performance.

At the same time, data centres are reaching a new scale. Driven by **artificial intelligence (AI), high-performance computing (HPC) and cloud expansion,** facilities are growing significantly in **size and power density.** Large campuses now frequently **exceed 50 MW,** while individual data halls operate at multi-megawatt levels.

At this scale, the way that electrical infrastructure is designed becomes a guiding factor for the entire data centre. Power architecture directly influences:

- The level of redundancy ensuring the **availability** target of 99,999% (5 nines)
- The capacity **usage rate**
- The **modularity** of the data centre for **future evolution**
- The **complexity of power distribution** to supply IT loads
- Capital and operational expenditure **(CAPEX and OPEX)**
- **Maintenance strategies** and operational constraints
- The site's overall **equipment footprint**
- **The global data centre carbon footprint**

As infrastructure grows, traditional architecture is being challenged to meet customer expectations in terms of availability while improving expenses and overall data centre operations.

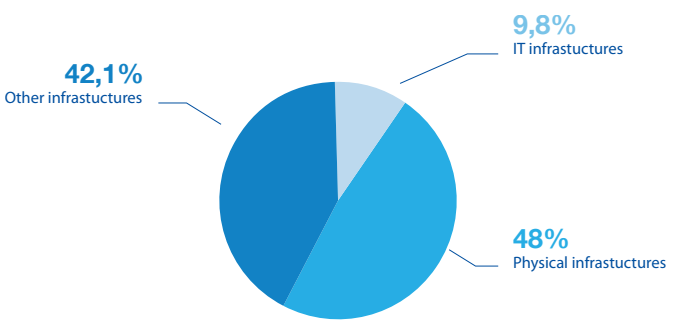
Why does architecture matter economically?

In colocation environments, physical infrastructure represents nearly half of total facility investment.

As power architecture shapes the entire electrical path, including transformers, generators, switchboards, UPS systems, batteries and electrical distribution, it directly influences the quantity of installed equipment and, consequently, the total cost of ownership (TCO) of the facility.

This makes power architecture a decision not only about resilience, but also about strategic investment.

Colocation data centre: Investment split by Infrastructure type



source : Frost & Sullivan 2025. Other infrastructures include aspects such as land, building, shell, etc.

Clément Barthelmebs
Data Centre Marketing Manager, SOCOMEC



The era of oversizing infrastructure for additional redundancy is over. Success now lies in delivering the right level of availability while maximising flexibility and infrastructure efficiency.”

Redundancy as an architectural choice

As facilities grow, architectural choices increasingly influence infrastructure efficiency, operational simplicity and long-term scalability.

Several power architecture models can achieve high availability objectives while structuring redundancy in different ways.

These approaches not only determine resilience levels, but also influence infrastructure utilisation, distribution complexity and future scalability.

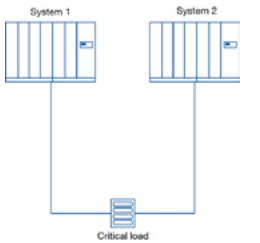
Redundancy should, therefore, be considered as a structuring design parameter, not only as a requirement in terms of reliability.

Common redundancy architectures include:

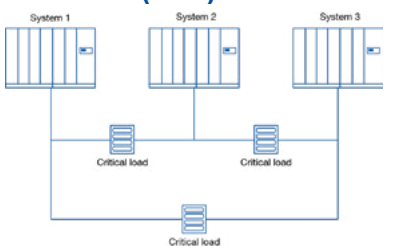
- **2N architectures, where the entire electrical chain is duplicated, creating two fully independent power paths.**
- **Distributed redundancy architectures (such as 4N/3), where reserve capacity is shared across multiple power modules and distribution paths.**
- **Block-redundant or N+1 configurations, where additional capacity is added at component or system level.**

All of these models meet recognised availability standards (such as Tier classifications), but they rely on fundamentally different ways of organising redundancy.

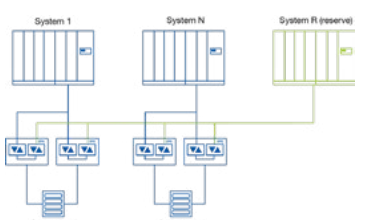
2N architectures



Distributed redundancy architectures (4N/3)

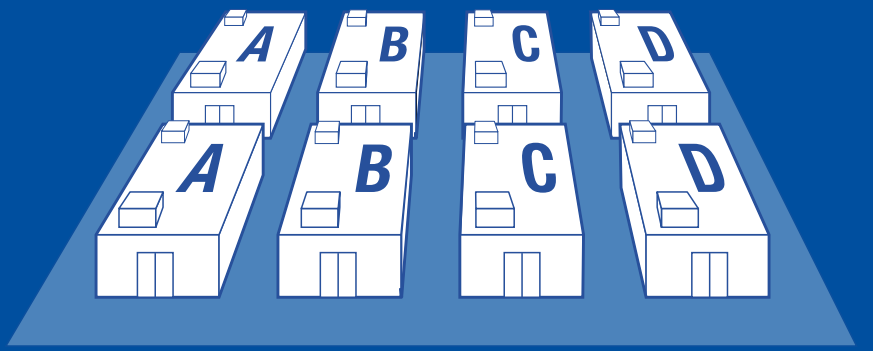


Block-redundant or N+1 configurations



Distributed redundancy architectural (4N/3)

The 4N/3 distributed redundancy architecture, often referred to as "4 makes 3", is widely used in large-scale data centres to achieve high fault tolerance while improving infrastructure utilisation.



Operating principle

In this architecture, the data centre is supplied by **four independent power feeds (A, B, C, D),** each backed by its own UPS system.

The system is designed so that the loss of any single UPS or power source can be absorbed by the three remaining systems.

This approach provides redundancy while avoiding the low utilisation typically associated with fully duplicated architectures.

Under normal operation:

- All four UPS systems share the IT load equally
- Each UPS typically operates up to **75% of its rated capacity**
- The load is distributed across all available power sources

In the event of a failure of one feed:

- The remaining three UPS systems automatically absorb the affected load
- Each UPS increases its output up to 100% capacity
- No interruption occurs to the IT load - Server PSUs automatically draw power from the remaining available sources

Distributed redundancy architectural (4N/3)

Key characteristics

High availability
No single point of failure.

No single point failure
4 system independant.

Distributed redundancy
Backup capacity shared across active power streams.

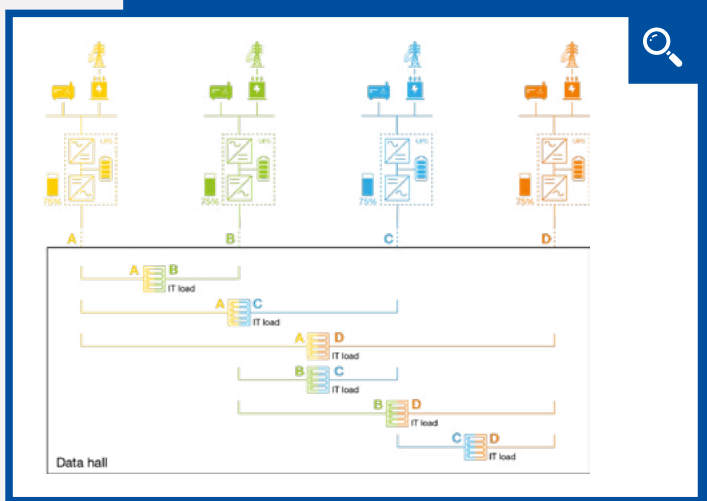
Distribution complexity
Multiple feeds and interconnections throughout the white space.

Operational management
Requires load balancing and infrastructure coordination.

Infrastructure utilisation
UPS systems typically operate at around 75% capacity under full load conditions.

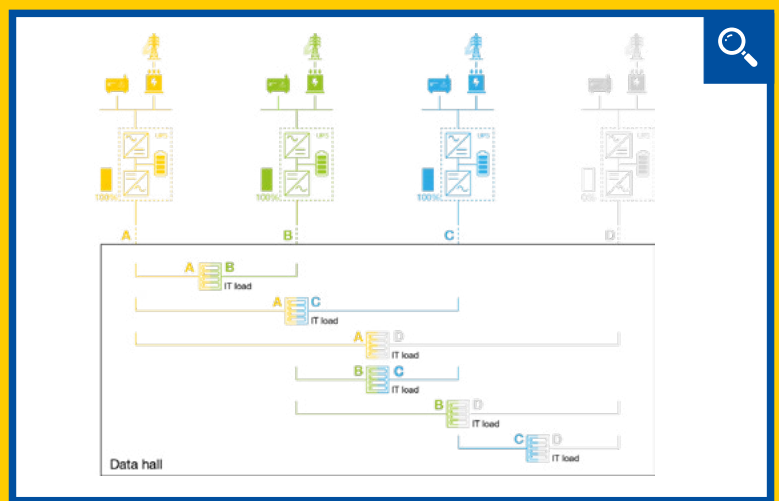
Normal operating mode

4 feeds | Up to **75%** usage rate



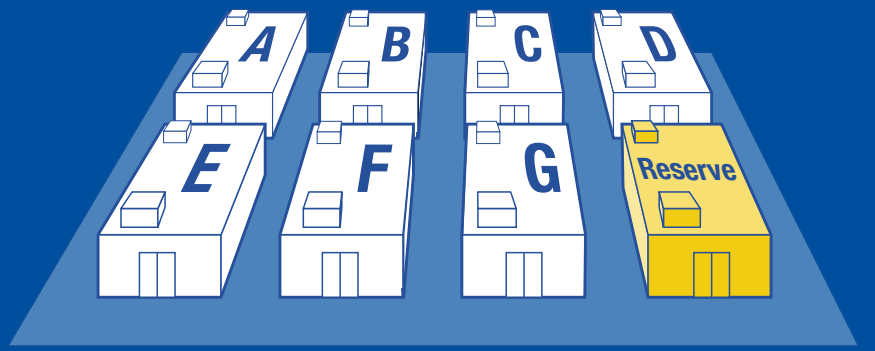
In the event of a failure or planned maintenance

3 remaining feeds | Up to **100%** usage rate



Block-redundant architecture (Catcher)

The Catcher architecture, also referred to as block-redundant architecture, centralises redundancy through a shared reserve source, maximizing the utilisation of active power infrastructure and reducing overall architecture costs.



Operating principle

In this architecture, the data center is divided into several independent primary power blocks supplying dedicated groups of IT loads.

In this configuration:

- N primary power sources feed the load,
- 1 redundant source is installed as a shared reserve for the entire group,
- each IT load is connected to both its primary source and the reserve via a Static Transfer Switch (STS).

Under normal conditions:

- Each primary line carries its assigned share of the load, typically close to full utilisation,
- The redundant line remains idle and unloaded,
- Each load is supplied exclusively by its primary source through the STS.

In the event of a failure or planned maintenance:

- The affected load is transferred to the reserve power block,
- The transfer is automatic without any wdisturbances,
- The remaining primary systems continue operating without redistribution.

This differs fundamentally from distributed architectures. Instead of spreading spare capacity across multiple active systems, the Catcher model relies on a standby source ready to replace any failed primary line.

Block-redundant architecture (Catcher)

Key characteristics

Shared redundancy
Backup capacity shared across multiple power blocks.

Simplified distribution
Fewer active feeders and interconnections throughout the white space.

Operational management
Simplified maintenance with fewer active components.

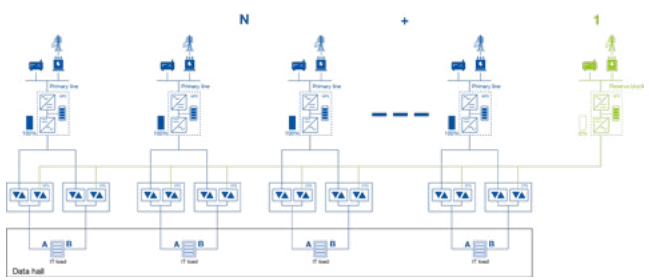
Infrastructure utilisation
Primary power blocks may operate close to full capacity under normal conditions.

Scalability & flexibility
Modular architecture enabling phased capacity expansion as demand grows

Sustainability
Fewer power assets and distribution lines reducing overall footprint

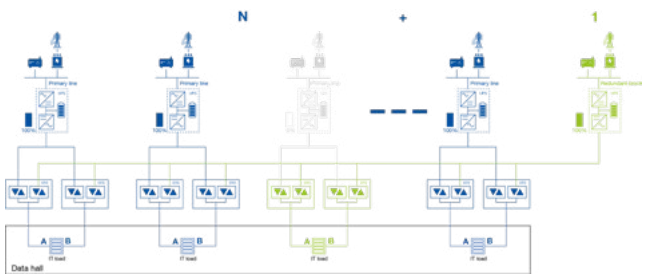
Normal operating mode

N + R feeds | Up to **100%** usage rate



In the event of a failure

(N-1) + R remaining Feeds | Up to **100%** usage rate



Block-redundant architecture (Catcher)

Several Catcher architectures can be implemented in a data centre. The main difference lies in the way the redundant block is connected and transferred to the A-side and B-side power inputs of the IT racks in the event of a power source failure.

Three typical architectures can be considered:

Dual STS Catcher architecture

A and B servers power paths are connected to the redundant block through dedicated STS units.

- In the event of a power source failure, both A-side and B-side streams are transferred immediately to the redundant block.
- Continuous power supply is maintained to all IT loads within the ITIC curve.
- Full dual-path redundancy is preserved.

Highest level of resiliency

SEE ARCHITECTURE 1

Single STS Catcher Architecture

One power path is connected to the redundant block through a dedicated STS.

- In the event of a power source failure, the protected power stream is transferred immediately to the redundant block.
- IT loads remain continuously supplied within the ITIC curve through one fully protected power path.
- One fully protected power path is maintained.

CAPEX-optimised while maintaining one fully redundant power path

SEE ARCHITECTURE 2

STS + ATS Catcher Architecture

The redundant block is connected through a combination of STS and ATS devices.

- In the event of a power source failure, one power path transfers immediately through the STS while the second power path transfers through the ATS (typically within 500 ms).
- Continuous supply is maintained to IT loads.
- Dual-path redundancy is preserved following the power source failure.

Balanced trade-off between resiliency and CAPEX

SEE ARCHITECTURE 3

Why is STS central to block-redundant architectures?

In a block-redundant architecture, the Static Transfer Switch (STS) is the component that connects the critical load to either the primary power source or the reserve power line.



Under normal operating conditions, the load is supplied through its primary path. In the event of a fault, maintenance operation or source degradation, the STS automatically transfers the load to the reserve path, ensuring continuity of supply.

The STS, therefore, does more than simply switch between sources: it enables redundancy to be structured at block level while maintaining the fault-tolerant behaviour required in mission-critical data centres.

Key functions of the Static Transfer Switch (STS)

Automatic source transfer
Transfers the load from the primary path to the reserve path in the event of a fault, maintenance operation or source degradation.

Fault isolation
Helps isolate electrical disturbances and prevents faults from propagating between power paths.

Extra embedded input breakers within the STS enhance system selectivity.

Short-circuit propagation prevention
STS intelligence detects abnormal conditions and prevents short-circuit faults from propagating across the redundancy block.

Intelligent load management
Embedded control functions manage transfer sequences to prevent overloading the reserve path during abnormal conditions.

Proven reliability
Modern STS technologies offer multi-million-hour MTBF and a proven track record in mission-critical environments.

Use case
60 MW IT data centre

To illustrate the structural differences between distributed and block-redundant architectures, a simplified use case can be considered at data hall scale.
This comparison highlights how the organisation of redundancy directly influences the amount of installed infrastructure required to support a given IT load, and why architecture becomes an increasingly important design parameter as facilities scale.



Use case definition

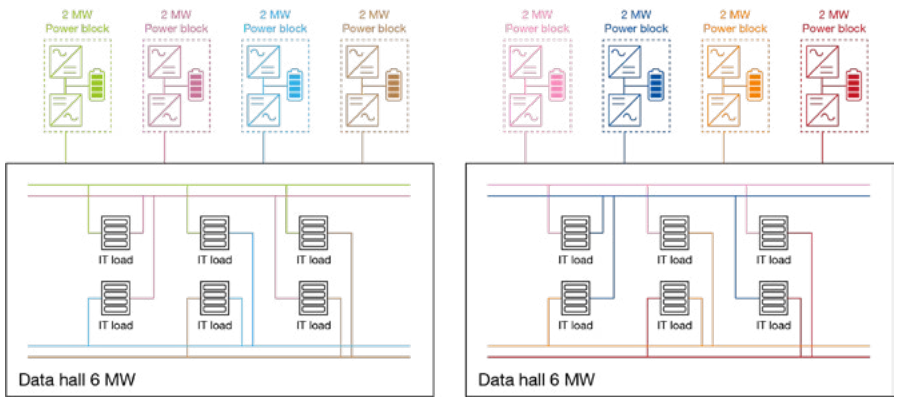
The reference scenario considers:

- 10 data halls each of 6 MW
- Total IT load: 60 MW

Tier 3 data centre design is considered as target using two different architectural approaches:

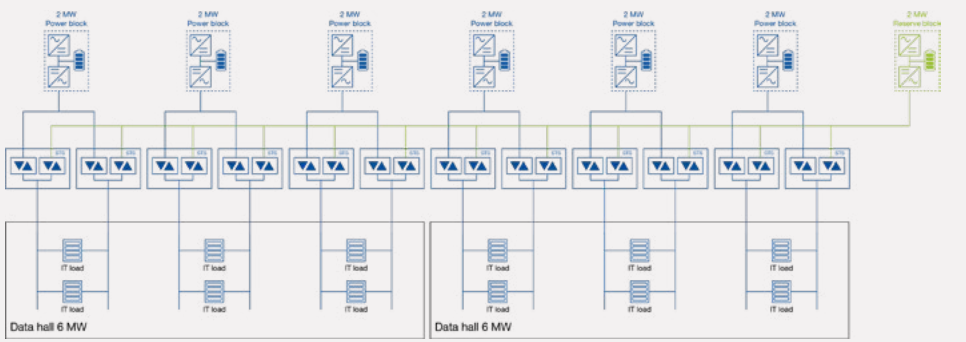
- a distributed architecture 4 makes 3 (4N/3)
- a block-redundant architecture (6+1)

Distributed architecture (4N/3)
12 MW IT load



SEE MORE

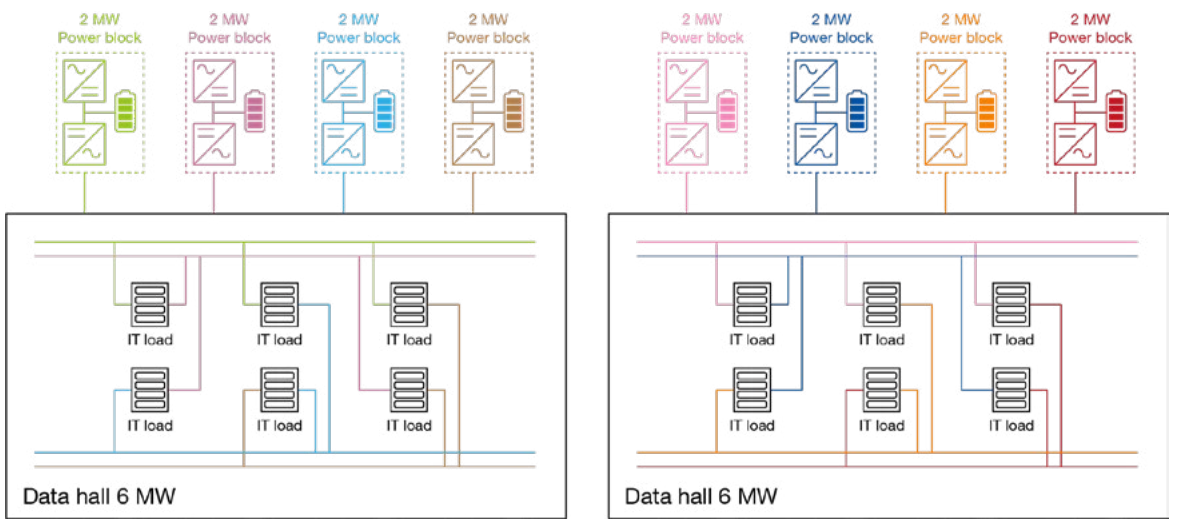
Block-redundant architecture (Catcher) 6+1
12 MW IT load



SEE MORE

Distributed architecture (4N/3)

BACK ARCHITECTURES



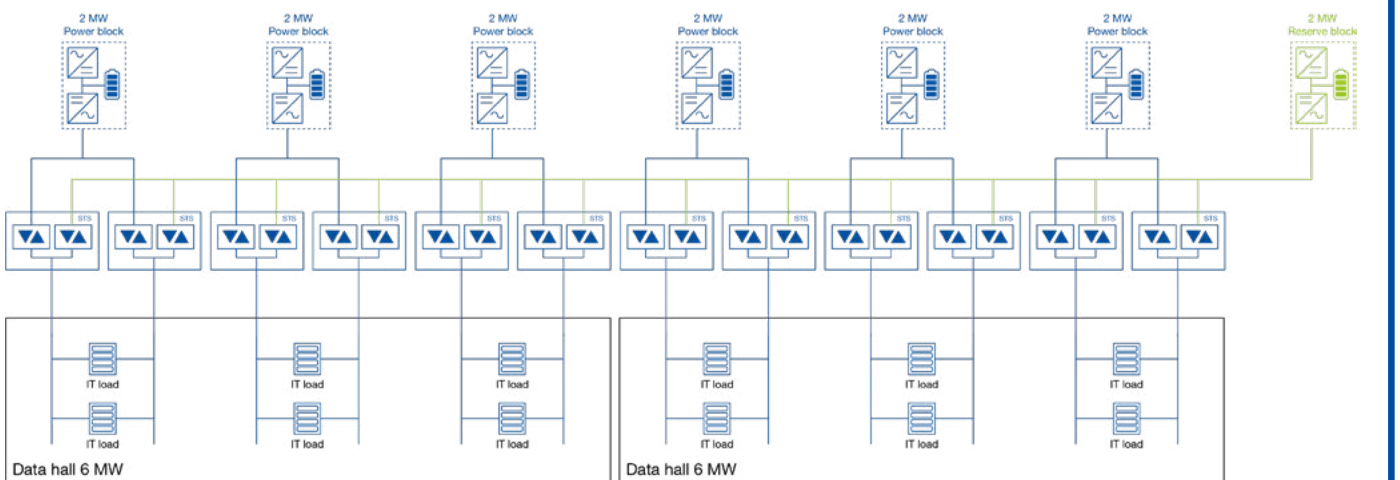
2 data halls represent

16MW total installed power for 12MW IT load

Meaning 25% trapped capacity not used most of the time

Block-redundant architecture (Catcher)

BACK ARCHITECTURES



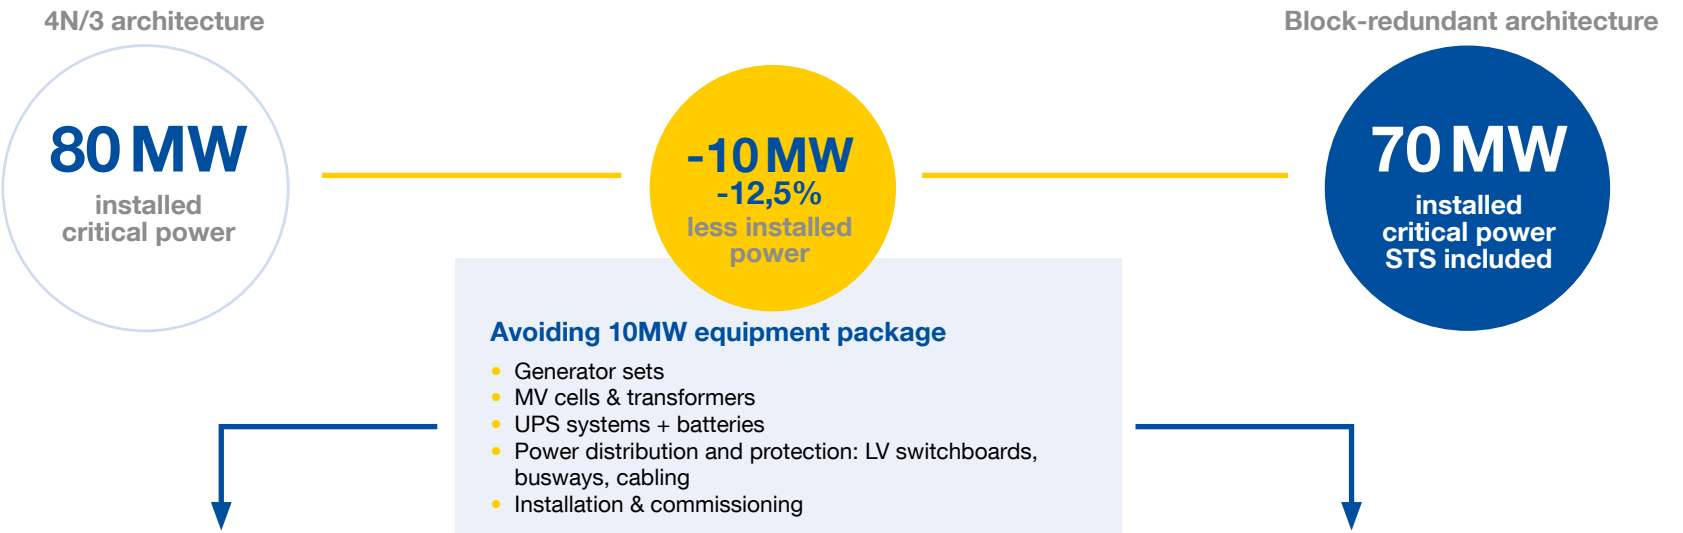
2 data halls represent

14MW total installed power for 12MW IT load

Meaning 1/6 trapped capacity not used most of the time

Same availability. 10 MW less infrastructure.

60 MW IT data centre – 4N/3 versus block-redundant



Avoiding 10MW equipment package

- Generator sets
- MV cells & transformers
- UPS systems + batteries
- Power distribution and protection: LV switchboards, busways, cabling
- Installation & commissioning

CAPEX

Millions of potential saving

Less equipment to purchase and commission

Footprint

> 80% net equipment-footprint saving 10 MW versus STS installed

Less equipment to manufacture, transport, install, lower embodied carbon footprint

Operations

Fewer assets, simpler operability

Less preventive maintenance, battery replacement, generators tests, …

Block redundancy turns infrastructure efficiency into lower TCO — without compromising availability.

Key takeaways

Availability remains non-negotiable

Different architectures can meet the same resilience requirements.

01 Architecture matters at scale

As facilities grow, architectural choices increasingly influence infrastructure investment (CAPEX), operational simplicity and long-term scalability.

02 Redundancy can be structured differently

Distributed and block-redundant architectures achieve availability objectives through different approaches.

03 Infrastructure efficiency becomes increasingly important

The impact of architectural choices extends beyond resilience to include footprint, maintenance requirements and infrastructure utilisation.

04 Block-redundant architectures provide an alternative approach

By sharing reserve capacity through STS-enabled transfer technology, block-redundant architectures offer a way to support high availability while reducing infrastructure duplication.

thank you SO much!

socomec
Innovative Power Solutions