



DATA CENTER FACILITY – LOW VOLTAGE

DIRECT CURRENT POWER DISTRIBUTION

VERSION 1.1.0

Date: 2 SEPTEMBER, 2026

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

Alliance Partners

Alliance Partner
Current/OS
Open Direct Current Alliance ODCA

Sub-Project Leads & Work Stream Leads:

Name	Company
Antonello Antoniazzi	ABB
Aurelio Rodriguez Echevarria	AMD
Brian Deckard	Siemens
Ehsan Nasr	Individual contributor
Jing Wang	Google
Joshua (JP) Buzzell (Sub-project Lead)	Individual contributor
Ke Sun	Vertiv
Laurens Mackay	DC Opportunities
Lisa Grams	UL Solutions
Paolo Catapane (Sub-project Lead)	ABB
Rudy Wang	Delta Electronics
Scott Hunter	UL Solutions
Tony Landry	Schneider Electric

Date: 2 SEPTEMBER, 2026

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

Contributors:

Name	Company
Abiezer Tejeda	TE Connectivity / Schaffner EMV
Adam Anders	Cyient Semiconductors
Adam Selvey	Jacobs
Al Warner	Dragon Q Energy
Allegia (Gia) Wiryawan	Schneider Electric
Aram Khodamoradi	Infineon
Arnaud Brozek	Socomec
Aurelio Rodriguez Echevarria	AMD
Ben Gully	SPOC Energy
Binesh Kumar	Atom Power
Bob Delisi	UL Solutions
Brian Zahnstecher	Nyobolt
Clayton Gibbons	SPOC Energy
Danish Zia	UL Solutions
Diego Gonzalez	Leibniz Institute for Plasma Science and Technology
Elijah Auger	Utilidata, Inc
Enrico Mazzola	TE Connectivity / Schaffner EMV
Eric Meier	Individual contributor
Erik Fosselmann	Danfoss
Franco Moriconi	VEIR Inc
Frederick P. Reyes	UL Solutions
George Zapf	Vertiv
Gerald Bode	nlevel

Date: 2 SEPTEMBER, 2026

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

Name	Company
Giel van der Broeck	Direct Energy Partners
Gonzalo Crespo Roche	DC Opportunities
Guang Yang	Siemens
Guido Macchi	Tokamak Energy
Haiguo Li	Delta Electronics
Haohao Yi	Individual contributor
Harry Soin	Advanced Energy
Harry Stockman	DC Expert
Hartwig Stammberger	Eaton
Hasala Dharmawardena	NERC
Jason Barmann	EPC Power
Jason Hopkins	UL Solutions
Jesse Echeverry	DC Opportunities
Jesus Sanchez Rodriguez	Vertiv
Jonas Knapp	Fraunhofer IPA
Justin Mayle	Vertiv
Keith Waters	Schneider Electric
Kimberly McGrath	Alderbuck Energy
Kiran Kutty	Heron Power Electronics
Liran Katzir	SolarEdge
Luiz Huet de Bacellar	Eaton
Marc Klimpel	Phoenix Contact
Marco Carminati	Schneider Electric
Marco Riva	ABB
Martin Ehlich	Lenze

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

Name	Company
Mike Jensen	UL Solutions
Mike Tu	NVIDIA
Misha Kumar	Delta Electronics
Naiqian Ren	Analog Devices
Nathan Atkinson	Nyobolt
Nick Powers	Infineon
Patrick Chambon	Schneider Electric
Pranesh Somasundaram	Amperesand
Randall Dollar	Siemens
Robert Osborne	UL Solutions
Rounak Siddaiah	Heron Power
Roy Ball	Mersen
Roy Davis	OmniOn Power
Saad Lambaz	Littelfuse, Inc.
Satarupa Bal	Google
Sebastian Breisch	ABB
Sergey Kynev	Siemens Energy
Seshadri Sivakumar	Individual contributor
Sina Driss	Jacobs
Siebe Berveling	Sensata
Somasundaram Essakiappan	Trane
Susie Choi	DG Matrix
Thomas Schuldt	DG Matrix
Thorsten Strassel	ABB
Todd Lottman	Eaton

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

Name	Company
Tom Domitrovich	Eaton
Tyler Andrews	Echion Tech
Uma Manickam-Chauhan	Siemens
Venkata Chivukula	Microsoft
Vibhor Iyer	Velankani Group
Vincent Sabalvaro, IV	UL Solutions
Wolfgang Voss	Siemens Energy
Xin Zhou	Eaton
Young Liang	XMight Smart Energy
Zafer Cankurtaran	Rittal

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

Version History

Date	Version #	Authors	Description
4 MARCH 2026	1.0	See previous page	First release
2 SEPTEMBER 2026	1.1.0	See previous page	Added new chapters 9, 10 and 11 (IT rack, Control&Operation, EMC), major expansion for all other chapters.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

Executive Summary

The rapid escalation of AI-driven workloads is pushing data center power systems beyond the limits of traditional architectures. Rack power densities are increasing sharply, load profiles are becoming more dynamic, and efficiency, scalability, and grid interaction are now first-order design constraints. To address these challenges, the data center industry is increasingly exploring direct current (DC) power distribution as a complementary, more power dense and energy efficient approach to conventional AC-based systems.

This white paper, developed within The Open Compute Project, a collaborative industry initiative focused on open, scalable, and efficient data center infrastructure, provides a high-level overview of DC power distribution architectures for data center facilities. It examines key design concepts, system boundaries, and tradeoffs across power conversion, grounding and earthing, protection, energy storage integration, and regulatory considerations.

Rather than promoting a single solution, the paper presents representative DC power distribution approaches in use or under evaluation across the industry. It is intended to support a shared understanding among data center operators, designers, equipment manufacturers, utilities, and Authorities Having Jurisdiction (AHJs), enabling informed discussion and consistent interpretation as DC architectures continue to evolve. This white paper represents an initial framework and identifies areas requiring further development.

Date: 2 SEPTEMBER, 2026

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

Table of Contents

Contents

1. Introduction....15
2. Compliance with Open Compute Project Tenets....16
2.1. Openness....16
2.2. Efficiency....16
2.3. Impact....16
2.4. Scale....17
2.5. Sustainability....17
3. Background Context....18
3.1. Data Center Transition to Direct Current: Background and Needs....18
3.2. Document Scope....21
4. System Architecture....23
4.1. Introduction....23
4.1.1. Definitions....23
4.2. From MV to LV - high level stages....25
4.2.1. Stage 1....26
4.2.2. Stage 2....30
4.2.3. Stage 3....31
4.3. Redundancy Levels....31
4.4. Power distribution redundancy schemes for DC data centers....34
4.4.1. Background....35
4.4.2. Reserve or catcher redundancy topology (N+x)....37
4.4.3. Dynamic bus topology (1.5N or 2N)....38
4.4.4. Dynamic redundancy topologies at rack level....44
4.4.5. Summary....48
4.5. Inclusion of redundant power generation....49
4.6. Verified Benefits of DC installations....49
5. Voltage Levels and Power Quality....51
5.1. Voltage Classes....51
5.2. Low Voltage (LV), Including Sub-Classes (e.g. ELV, LES)....51
5.2.1. General Definition....51
5.2.2. Extra-Low Voltage (ELV)....53
5.2.3. Limited Energy Systems (LES)....54
5.3. High Voltage (HV), Including Sub-Classes (e.g. MV, HV, EHV, UHV)....55
5.3.1. General Definition....55
5.3.2. Medium Voltage (MV)....55
5.3.3. High Voltage (HV), Extra High Voltage (EHV), Ultra High Voltage (UHV)....56

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

5.4. Nominal Voltage....56
5.5. Rated Voltage vs. Operating Voltage....57
5.5.1. Slash Ratings, Neutral, and Midpoint....59
5.5.2. Voltage Insulation and Isolation....60
5.6. Voltage Bands in Active DC Microgrids....61
5.6.1. Voltage Band Definitions....62
5.6.2. Existing Voltage Bands....63
5.6.3. Narrow vs. Wide Voltage Bands....66
5.7. Interoperable Voltage Band Proposal....69
5.7.1. Low Voltage Limit Considerations....74
5.7.2. Voltage Band Impact of Static Line Losses....75
5.8. Power Quality and Grid Interaction....77
5.8.1. AC Power Quality at the Utility Interface....77
5.8.2. DC Power Management with Droop Curves....78
6. Polarities & System Earthing....80
6.1. Polarities....80
6.1.1. Bipolar....82
6.1.1.1. Bipolar with Distributed M....83
6.1.1.2. Bipolar with Partially Distributed M....84
6.1.1.3. Bipolar without Distributed M....85
6.1.2. Unipolar....86
6.2. System Earthing....87
6.2.1. Isolation of Electrical Systems....89
6.2.1.1. Medium Voltage (MV) to Low Voltage (LV)....90
6.2.1.2. Low Voltage (LV) to Extra Low Voltage (ELV)....90
6.2.1.3. Grounded AC side Low Voltage AC (LVAC) to Low Voltage DC (LVDC)....91
6.2.2. TN-S (Solidly Grounded System)....92
6.2.2.1. Direct Touch Protection (with RCD)....93
6.2.2.2. Multiple References to Earth....93
6.2.3. IT (Ungrounded System or High Impedance Grounded)....95
6.2.3.1. Direct Touch Protection (with IMD)....95
6.2.3.2. Y-Capacitor Limitations....96
6.2.4. Corrosion....97
6.3. Combinations of Polarities, System Earthing and Voltages....98
6.3.1. Combination Table Analysis: Overvoltage Category (OVC), Impulse Withstand Voltage (Uw, Uimp), Clearance, and Creepage....101
6.3.1.1. Additional Discussion: Relationship between product standards and electrical spacing for NEC, UL, and IEC....102
6.3.1.2. Additional Discussion: Isolation requirements in system design....103
6.3.1.3. Additional Discussion: Voltage with respect to earth for different systems

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

(polarities)..... 104
6.3.1.4. Additional Discussion: Line-to-line clearance..... 104
6.3.2. Combination Table Analysis: Device Compatibility (350V or 400V device compatible, 700V or 800V device compatible, 1400V or 1500V compatible)..... 104
6.3.3. Combination Table Analysis: Wiring (Active Wires, Active Wire Material, Relative Cable Losses, Relative Current for same balanced Power)..... 105
6.3.4. Combination Table Analysis: Impact on Protection Devices..... 106
6.4. Conclusion..... 107
6.4.1. Likely to be used Combinations..... 108
6.4.2. Less likely to be used Combinations..... 109
7. Protection and Safety..... 110
7.1. Introduction..... 110
7.2. Protection against overcurrent..... 111
7.2.1. Overcurrent faults in active DC systems..... 111
7.2.2. Selectivity..... 121
7.2.3. Branch vs supplementary protection..... 123
7.2.4. Dependence of protection devices on grounding..... 123
7.2.5. Overcurrent protection devices..... 124
7.2.5.1. Fuses..... 125
7.2.5.2. Electromechanical Circuit breakers..... 125
7.2.5.3. Solid state Circuit breakers(SSCBs), Semiconductor Hybrid Circuit breakers(SHCBs) and e-fuses..... 126
7.3. Protection against electric shock..... 129
7.4. Protection against thermal effects and fires..... 131
7.5. Protection against voltage disturbances..... 132
7.6. Protection against arc-flash..... 133
8. System Main Components..... 136
8.1. Utility Grid..... 137
8.1.1. Voltage Ride Through (VRT)..... 138
8.1.2. Frequency Ride-Through (FRT)..... 140
8.1.3. Post Fault Active Power Recovery (PFAPR)..... 141
8.1.4. Protection System Coordination..... 143
8.1.5. Ramp Management..... 143
8.1.6. Oscillation Management..... 143
8.1.7. Power Quality..... 143
8.1.8. Monitoring Methods..... 143
8.1.9. Reactive Power and Voltage control..... 143
8.1.10. Validated Grid Models..... 144
8.2. On-Site Generation..... 144
8.3. AC/DC Conversion..... 144

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

8.4. ESS....151
8.4.1 Technology....153
8.4.2 ESS Safety....157
8.5. Distribution....159
8.6. AI Rack....164
8.7. Miscellaneous....165
9. Interface with Rack & Power....166
9.1. Components of an LVDC Input Rack....167
9.1.1. Junction Tray/Box....167
9.1.2. LVDC Vertical Bus bar....167
9.1.3. LVDC Smart Rack PDU....167
9.1.4. LVDC Input Conversion Shelf....168
9.1.5. Rack Architecture With an LVDC to ELVDC Conversion Shelf....169
9.2. Rack Architecture with an LVDC to ELVDC Conversion Shelf....170
9.3. IT Rack With LVDC-native IT Tray....171
9.4. Preliminary Rack Level Voltage Specifications....172
9.4.1. Input Voltage Range....173
9.4.2. Continuous Operation Definition....173
9.4.3. Transient Level 1 Operation Definition....174
9.4.4. Transient Level 2 Operation Definition....174
9.4.5. Transient Level 3 Operation Definition....174
9.5. LVDC Input Power Supply Architecture and Components....175
9.5.1. Overview....175
9.5.2. IT Tray Input Connector....175
9.5.3. Hot Swap Controller Circuitry....176
9.5.3.1. Initial Inrush....177
9.5.3.2. Charging Current and Startup Sequence....178
9.5.3.3. Monitoring and Telemetry....179
9.5.3.4. Discharge....180
9.5.3.5. Internal Auxiliary Power Supply....180
9.5.4. DC/DC Converters (LVDC/ELVDC)....180
9.5.4.1. Connection Recommendations....180
9.5.5. Isolation....181
9.5.6. IT Tray Grounding....182
9.6. Energy Storage and Power Oscillation Mitigation....182
9.6.1. Rack Level Energy Storage To Handle The Pulsating Power....183
9.6.2. Distribution Level Energy Storage To Meet The Grid Requirements....185
9.6.3. Hierarchical Energy Storage Architecture....186
10. System Operation, Control and Monitoring....188
10.1. MOP and SOO in Traditional AC Data Center Architectures....188

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

10.2. Transitioning to the DC Paradigm: From Human-centric MOP to System-centric SOO... 188
10.3. The "Wild West" of LVDC Native Datacenter Control: Undefined Boundaries.....189
10.4. SOO Scope for an MVAC-to-LVDC Data Center..... 189
10.5. Architectural Control Boundaries & Interoperability..... 190
10.5.1. Control boundaries.....190
10.5.2. Interoperability and Communications Protocols.....191
10.5.3. Protocol Domains.....192
10.5.4. Communications, Cybersecurity, and Event Evidence..... 193
10.6. Deterministic Sequence of Operations (SOO).....194
10.6.1. System Operation, Control and Monitoring Aspects..... 194
10.6.2. MVAC Source Acceptance and LVDC Bus Energization..... 194
10.6.3. Startup Sequence..... 194
10.6.4. Black Start.....194
10.6.5. Precharging.....194
10.6.6. Power Management with Droop Curves..... 195
10.6.7. Rack Load Admission, IT Coordination, BESS, and Cooling Interlocks..... 196
10.6.8. Utility Interface, Ride-Through, and Grid-Disturbance SOO.....196
10.6.9. DC Bus Sectionalization, Protection, and Fault Isolation.....198
10.6.10. Thresholds and Failure Timing Budget..... 199
10.7. Grounding, Insulation Monitoring, and Personnel Safety Interfaces.....199
10.8. Stored Energy, Active Discharge, and Verified-Zero-Energy.....200
10.9. Minimum SOO Deliverables..... 200
10.10. Enabling Technologies for Microsecond Survival-New Equipments.....201
10.10.1. DC Fault Transient Recorder.....201
10.10.2. DC Sequence Controller (SOO Automation Unit).....202
11. Electromagnetic Compatibility (EMC).....203
11.1. Scope and System Context..... 203
11.1.1. AC-DC Power Converters.....204
11.1.2. DC-DC Converters.....204
11.2. Emissions Requirements..... 205
11.2.1. Applicable Standards..... 205
11.2.2. Conducted Emissions..... 205
11.2.3. Radiated Emissions..... 206
11.2.4. Harmonics and Ripple.....209
11.3. Immunity Requirements..... 209
11.3.1. Applicable Standards..... 209
11.3.2. Key Immunity Phenomena.....209
11.3.3. DC Specific Immunity Gaps..... 210
11.4. Grounding and Bonding Strategies..... 211

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

11.4.1. Grounding Topologies..... 211
11.4.2. Recommended Grounding Approach.....211
11.4.3. Bonding and Return Path Design.....212
11.5. Cable Shielding and Filtering.....212
11.5.1. Cable Shielding.....212
11.5.2. Filtering..... 212
11.6. System-Level EMC Considerations..... 213
11.6.1. Loading Considerations..... 213
11.6.2. Global Compliance.....213
11.7. Gaps and Future Work.....213
12. Codes and Standards.....215
12.1. Why DC Codes & Standards Matter..... 215
12.2. Key Challenges in DC Standardization.....215
12.3. System-Level Standards.....217
12.3.1. Existing system-level knowledge and applications..... 217
12.3.2. Existing system-level standards.....218
12.3.3. System-level gaps and standards in development..... 219
12.4. Product-Level Standards: Existing & In-Development.....220
12.4.1. Existing product-level standards.....221
12.4.2. Product-level standards in development.....222
12.5. Installation Codes: Existing & Gaps.....225
12.5.1. Global Installation Requirements.....226
12.5.2. United States Installation Requirements.....227
12.5.3. Emerging trends and recommendations.....229
12.6. Interim Guidance..... 229
13. The path forward - Key Takeaways for Stakeholders.....231
14. Conclusion..... 233
15. References..... 234
15.1. Papers..... 234
15.2. Web Resources..... 235
16. License..... 238
16.1. Creative Commons..... 238
17. About Open Compute Foundation..... 239

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

1. Introduction

As data center power systems evolve to support increasingly demanding compute environments, there is growing interest in alternative power distribution approaches that can accommodate higher power densities, dynamic load behavior, and expanding system scale. Among these approaches, direct current (DC) power distribution has emerged as an area of active evaluation and early deployment across the industry.

DC power architectures introduce different design considerations than traditional AC systems, particularly in areas such as power conversion boundaries, grounding and earthing, protection and fault behavior, integration of energy storage, and interaction with upstream utility infrastructure. These considerations span multiple layers of the data center, from facility-level power distribution to zone- and rack-level implementations, and involve close coordination between equipment design, system architecture, and regulatory requirements.

This paper focuses on describing these considerations at a conceptual and architectural level. It outlines commonly discussed DC power distribution configurations, identifies key system components and interfaces, and highlights technical tradeoffs that may influence design decisions. The intent is to establish a common technical framework rather than to recommend or mandate a specific topology. However, this white paper represents an initial framework. Areas requiring further development are identified within the document.

The content is written for a broad audience, including data center operators, designers, equipment manufacturers, utilities, and Authorities Having Jurisdiction (AHJs). All architectures discussed remain subject to applicable codes, standards, equipment certification, and AHJ approval. As data center power systems continue to evolve, multiple AC and DC distribution models are expected to coexist, shaped by site-specific requirements, regional practices, and ecosystem maturity.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

2. Compliance with Open Compute Project Tenets

This white paper aligns with the core tenets of Open Compute Project by promoting openness, efficiency, impact, scalability, and sustainability in the context of data center DC power distribution. As a conceptual and architectural document rather than a hardware or software contribution, its primary objective is to enable shared understanding, informed discussion, and ecosystem alignment across a diverse set of stakeholders.

2.1. Openness

This paper supports openness by presenting DC power distribution concepts, architectures, and design considerations in a vendor-neutral and non-proprietary manner. It avoids prescribing specific products or implementations and instead focuses on commonly discussed approaches, tradeoffs, and system-level implications. By establishing a shared technical vocabulary and framing considerations applicable across regions and deployment models, the document enables broad participation, review, and extension by the wider data center ecosystem.

2.2. Efficiency

Efficiency is addressed through a system-level examination of DC power distribution, including the potential reduction of power conversion stages, improved utilization of conductors, and closer integration of energy storage. The paper emphasizes architectural choices that can influence electrical efficiency, power density, and operational effectiveness, while recognizing that realized gains depend on site-specific design, load profiles, and implementation details.

2.3. Impact

The document aims to create industry impact by clarifying emerging DC power distribution concepts that are increasingly relevant to next-generation, high-density data centers. By providing a high-level, accessible reference for designers, operators, equipment manufacturers, utilities, and Authorities Having Jurisdiction (AHJs), the paper supports more consistent interpretation, reduces ambiguity, and facilitates informed evaluation of DC architectures across the ecosystem.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

2.4. Scale

Scalability is addressed by considering DC power distribution across multiple layers of the data center, including facility, zone, and rack-level implementations. The paper highlights how different architectural approaches may support incremental deployment, modular growth, and adaptation to increasing power densities. This perspective aligns with OCP's emphasis on designs that can scale efficiently from smaller deployments to large, hyperscale environments.

2.5. Sustainability

Sustainability is reflected in the discussion of architectural options that may reduce energy losses, enable higher system efficiency, and support integration with energy storage and advanced grid interaction strategies. By focusing on long-term infrastructure adaptability and efficient power delivery, the paper contributes to broader sustainability goals without promoting specific technologies or solutions as universally optimal.

The OCP Data Center Facilities (DCF) Sustainability Project was formed to align the DCF project with the fifth OCP tenet - Sustainability. OCP has an IT focused Sustainability Project that focuses on the decarbonization of the IT infrastructure. This DCF-Sustainability Project is focused on decarbonizing the non-IT aspects of the data center. This project will engage industry partners to develop technology roadmaps, standardized carbon/water characterization metrics and methodologies, and provide guidelines for the data center industry in support of achieving net zero carbon goals through innovation and collaboration in products and technologies. OCPs unique combination of hyperscalar influence, vendors and sub-projects such as the OCP Ready Program can be leveraged to accelerate the industry's sustainability goals.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

3. Background Context

3.1. Data Center Transition to Direct Current: Background and Needs

As AI models scale in complexity, significant architectural changes are envisioned to the IT rack to support higher compute density and speed up the overall compute cluster by minimizing network packet loss (through staying within the copper layer boundary and reducing reliance on optical interconnections). Low Voltage Direct Current (LVDC) Power Distribution (≤1500 VDC) will soon function as a critical link between the IT rack AI factory compute load and incoming data center facility power – whether from the utility service, local production, or both. To address this critical and rapidly developing area, The Open Compute Project (OCP) has organized a Power Distribution sub-project within the Data Center Facility project.

Initial adoption of the Mt. Diablo Power Rack (side car) enables the ORv3 HPR IT Rack to increase from AC-supplied HPR v1 33kW or HPR v2 72kW power shelves to DC-supplied HPR v4 100kW power shelves within the same 1 OU space footprint – itself a 39% increase in power density per power shelf, made possible in part by removal of the AC-DC conversion stage and reduction in the bulk capacitance required for hold-up time.

Even if AC-supplied power shelves become more power dense, the constraint of the number of feeds going to the compute rack comes into play. Currently, the maximum AC feed size typically used by non-specialized personnel is 100A. Assuming a realistic maximum of eight 100A feeds de-rated to 80A in a single wide compute rack at 415 VAC would yield a maximum rack density of ~460kW (no redundancy). A practical limit based on reuse of existing busway rated to 600 VAC would yield ~665 kW.

Additionally, moving in-rack energy storage (in-rack 2 OU BBU and optional CBUs) from the IT Rack to the Power Rack has the further advantage of freeing up approximately 12-16 OU space footprint within systems that use in-rack storage, yielding a further 27-36% increase in available 44 OU IT Rack space for further compute power. Furthermore, moving to DC distribution allows power rack or in-row energy storage to operate at higher voltage levels with greater energy density than existing 50 VDC storage within the IT rack, while supporting high utilization swings without the round-trip power loss penalty of unnecessary AC-DC conversion.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

Ultimately, achieving the roofline power potential of HPR v4 requires addressing the further bottleneck of existing 50 VDC vertical bus distribution within the IT rack (as even ORW racks with horizontal bus face power limitations beyond the 100's of kW). For this reason, higher voltage (e.g. 800 VDC) vertical bus represents the next progression for the IT rack, enabling it to realize the potential of its 1+ MW roof line (see Figure 3.1).

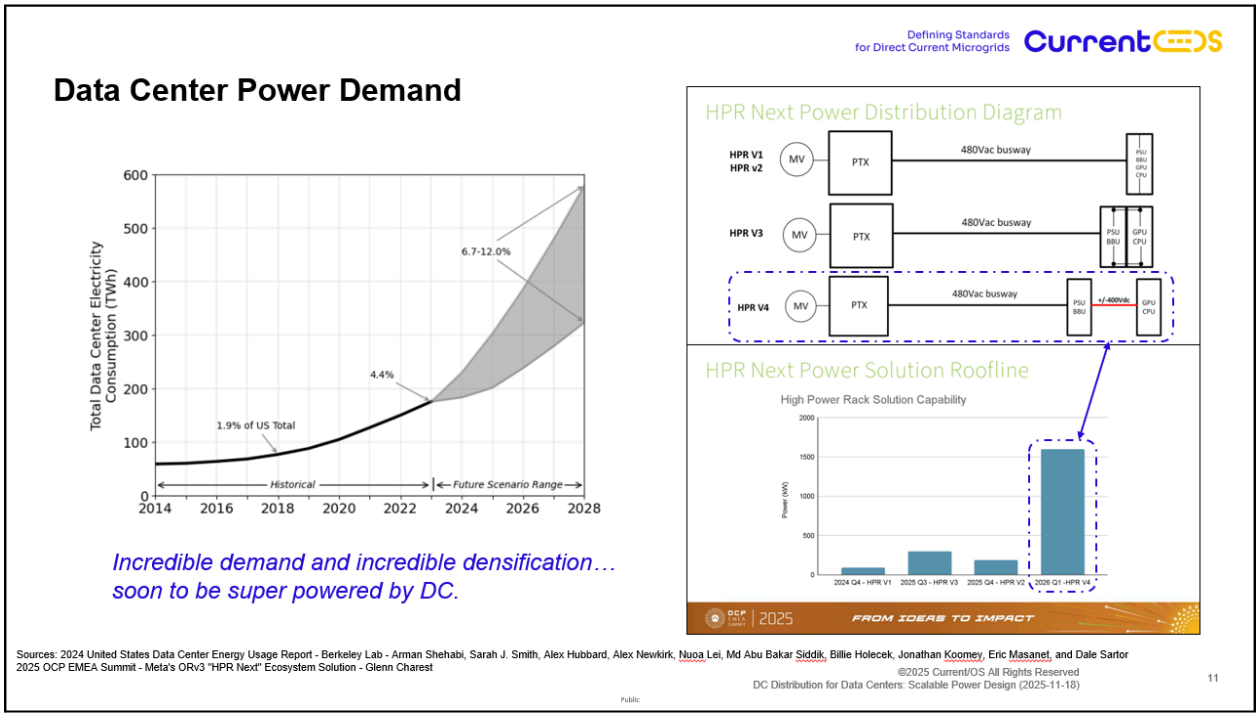


Figure 3.1. IT Rack Progression (Ref.: Current/OS: DC Distribution for Data Centers)

This progression could be realized by shifting from utilizing DC power shelves with DC PSUs for conversion to 50 VDC vertical distribution (typically using Front-End Power DC-DC modules) to utilizing DC junction boxes (capable of interfacing the same higher input DC voltage for vertical distribution) (see Figure 3.2 and Figure 3.3).

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

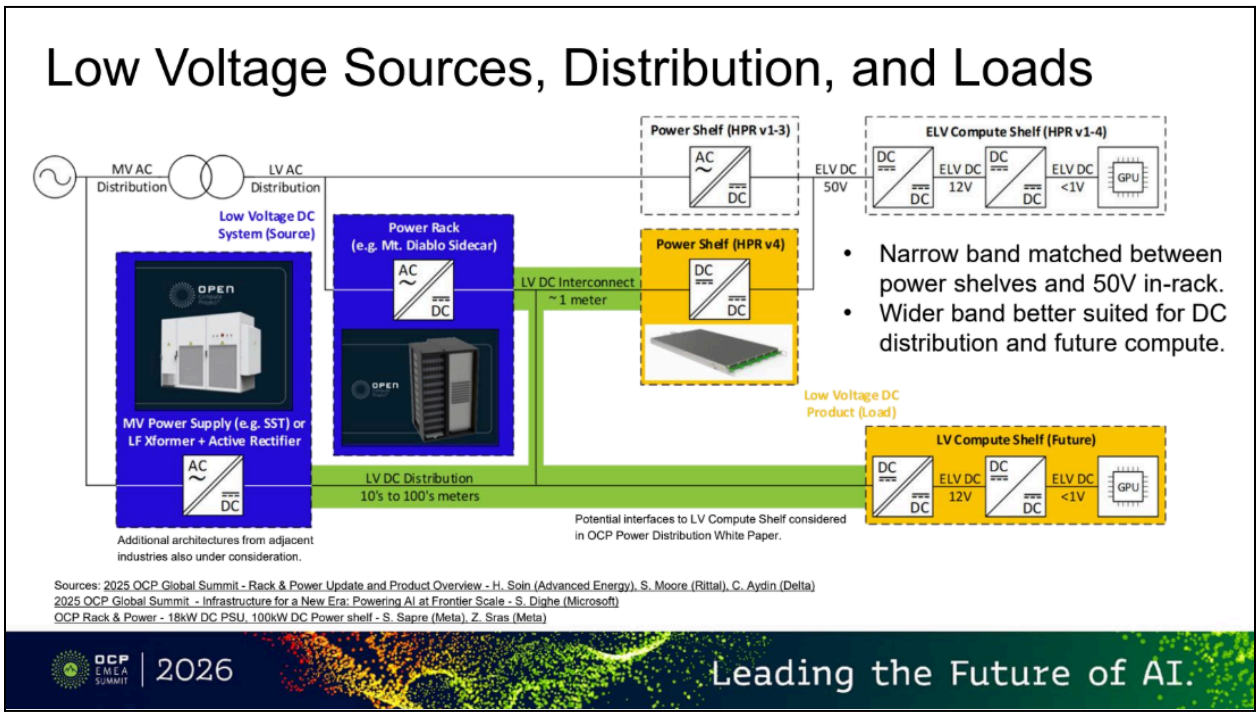


Figure 3.2. DC Distribution Progression (Ref.: Current/OS: Insights from OCP EMEA 2026)

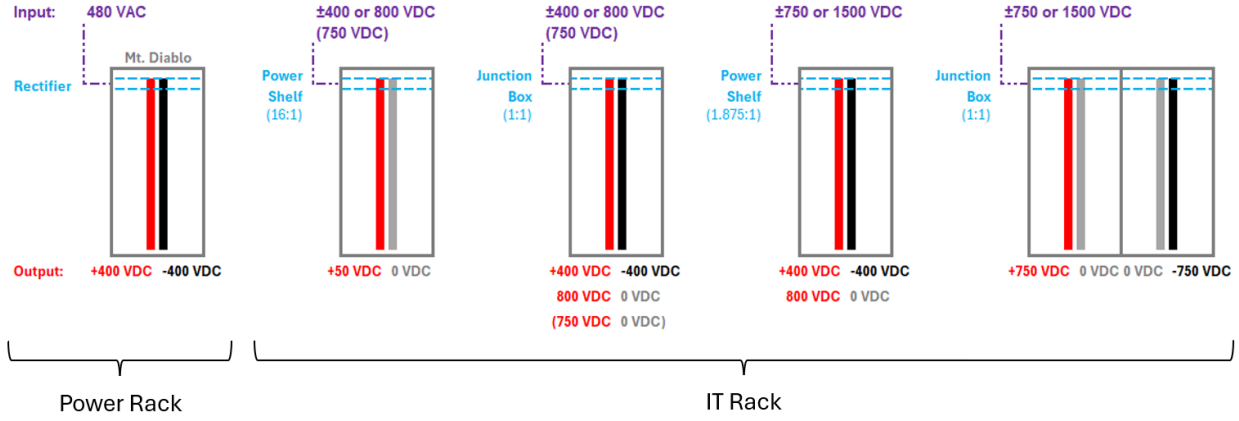


Figure 3.3. Example Interfacing to Power Rack and IT Rack DC Vertical Bus

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

3.2. Document Scope

This white paper is intended to address LVDC power distribution and its role as an interface between incoming Medium Voltage Alternating Current (MVAC) power (presently in the gray space) with IT rack compute in the white space. Section 4 and Section 8 of this white paper are intended to focus on the conversion of MVAC to LVDC, whether by MV Power Supply (commonly referred to as Solid State Transformer, SST) or other conversion components. Future editions of the white paper may also consider other possible LVDC sources, their interconnection, and varying levels of available fault current.

While energy storage is also included, direct connection to LVDC power distribution is excluded (i.e. a converter interface is required). It is important to note that upstream of the converter interface, energy storage presents significantly higher available fault current with additional requirements (outside the primary focus of this document) regarding circuit protection, restricted access, etc. according to equipment listing, manufacturer's instructions, and regulatory codes.

At present, the scope is limited to 1500 VDC, thereby excluding what is commonly referred to as Medium Voltage Direct Current (MVDC) which finds usage in other applications such as 2000 VDC for photovoltaic (PV) renewable energy sources. Future editions of the white paper may also consider utilization of an MVDC underlay grid and SST conversion (or otherwise) from MVDC to LVDC.

Additionally, while direct current usage for industrial / commercial building microgrids is typically focused on loosely regulated voltage droop control for power balancing of sources and loads (refer to Section 5.8.1) and interoperable component operation across a wide band of continuous voltage (reducing round trip losses for energy storage through the usable state-of-charge range), Direct Current usage in data centers is focused on tightly regulated operation in a narrow band of continuous voltage optimized for maximum power delivery at the limit of component voltage class capability.

These differences in application lead to a further consideration regarding the future trajectory of direct current usage in the data center facility for cooling, lighting, and other non-regenerative loads which have the potential to take advantage of efficiency from higher DC voltage distribution and reduced AC conversion losses. It is likely that such usage will reflect the higher continuous voltage limits of IT rack compute loads rather than the traditional limits of DC

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

microgrids. However, it may also introduce additional changes regarding bidirectional power flow considering energy storage, and is thus not included in the present scope.

Further, in the transition to DC-supplied IT racks, the power rack (sidecar) has been defined by the Mt. Diablo specification as developed by the OCP Rack & Power project. The scope of this white paper is limited to interface with DC-supplied IT racks including power shelves.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

4. System Architecture

4.1. Introduction

This subsection describes, at a high level, the connection between loads and power distribution. In order to simplify the discussion we consider two topologies: a single “line-up” for an MV to LV power block to supply the AI load (e.g. 3 MW+), and a direct current power generation source feeding a similar AI load / Data Center room. The focus of this paper is to be a foundational document to be built upon with further work.

4.1.1. Definitions

In the rest of this subsection the following terms will be used:

Redundancy

Redundancy is the practice of having backup equipment, power feeds, and cooling systems in place so that if one component fails, operations continue uninterrupted.

- N: The bare minimum infrastructure required to run at full capacity. There are no backups.
- N+1: For every critical component required, there is exactly one backup unit ready to take over.
- N+x: For every critical component required, there is exactly x amount of backups ready to take over
- 2N: Fully mirrored systems. There are two entirely independent sets of equipment capable of bearing the full load.

Note: In some cases the redundancy is described as the amount of critical components needed to run the system. This way the description is e.g. N-x or N/2. However, We have chosen the earlier mentioned nomenclature for this white paper.

Reliability

Reliability refers to the statistical probability of a system to perform its required functions at or during a specific time under specific conditions. When viewed in terms of time, this is called uptime.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

Failure Rate and MTBF (mean time between failures)

Failure rate is the frequency with which an engineered system or component fails within a certain period and is the inverse of MTBF. E.g. 0.6% failures per year means that on average 6 out of 1000 components fail each year, or that each component is expected to have an MTBF of 166.6 years. While reliability is the probability that a system will not fail, the failure rate focuses strictly on the statistical occurrence of breakdowns.

MTTR (mean time to repair)

MTTR is the average time required to troubleshoot, fix, and fully restore a failed system or component to operational status. This includes switching time to reestablish the connection, e.g. recloser operations.

Uptime or availability

Uptime is the percentage of time that a data center facility, network, or server remains fully operational and accessible over a specific period, typically a 12-month interval. For example, 99.995% uptime translates to a maximum of just 26.3 minutes of allowed downtime per year. A systems' availability is given in simplified form by:

$$A = MTBF/(MTBF + MTTR)$$

In a redundant system N+x, the system uptime will only be affected when x+1 items fail. If $MTTR_u$ is the MTTR value of the individual unit, the value of the MTTR system value in the above equation $MTTR_u/(x+1)$

Dynamic bus topology

Topology with critical components able to cover itself and half or a full other critical component, often referred to as 1.5N or 2N

Reserve or Catcher bus redundancy

Topology where critical components are backed up by x amount of units which are idle during normal operation. The reserve components are energized and ready but do not share the load in normal circumstances. These are referred to as N+x

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

Dynamic (reserve) redundancy topology

Topology where critical components are backed up by 1 or 2 units which are continuously engaged, thus not idle during normal operation. The reserve components share the load in normal circumstances. They are referred to as “N+x to make N”, e.g. 4 to make 3.

4.2. From MV to LV - high level stages

MV gets distributed from substation into smaller power nodes that feed into the data hall’s and mechanical hall’s LV system. These power nodes consist typically in a MV switchgear setup that allows a redundant routing of the power, e.g. a loop feed, that ties in several nodes (connection points) together and is fed from two different sources. Each node will convert the power needed into LV power being used at a load nearby via a MV/LV transformer. These nodes are not placed in parallel, but can be interlaced on the LV side with a reserve bus topology (for N+1 to make N) or dynamic bus topology (2N) to provide redundancy. One node can have several feeds, with each feed having a dedicated breaker and protection setup.

The current standard topology consists in only AC type distribution, with eventual conversion from AC to DC in the rack to DC infrastructure located downstream from the LV AC switchgear, for each dedicated feed. Building on this topology, a first step into conversion into a DC network could be the addition of a DC network at the LV root of the power node (e.g. as described by the Mt. Diablo specification). One transformer would be feeding at the same time AC and DC networks.

Transforming a data center’s power supply is more than a technical upgrade – it is an adoption of a staged approach and a strategic step toward securing performance, efficiency, and reliability. The first principle driver is to support compute densification and congestion management.

Such a transformation is seen to occur in adoption stages. This requires a structured, staged approach that brings clarity to complexity and stability to change–focusing on the first principle driver of the change, then moving to further optimizations as the supply chain matures to support size and scale.

Each phase provides valuable insights that help us optimize performance and enhance the system along the way. Based on this, we have structured the DC power transition into three distinct stages.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

4.2.1. Stage 1

- IT Racks (and optionally associated cooling) are supplied by direct current with other equipment powered by AC
- Option 1a: with isolation on the DC/DC converters for the rack, i.e., where the step-down conversion from 800 V to below 50 V happens; see Figure 4.1.

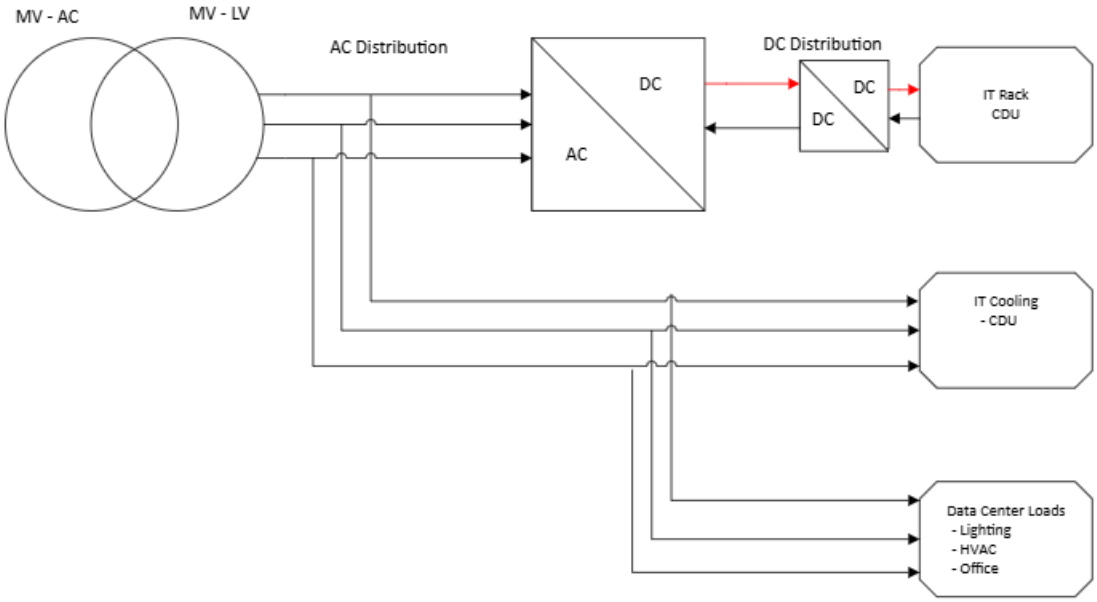


Figure 4.1. Isolation at the DC/DC converters for the rack

- Option 1b: An alternative for the same stage with isolation at the central AC/DC converter is shown in Figure 4.2 Isolation at the central AC/DC converter. This isolation allows for a different grounding system on the DC side from the AC side - more on grounding in chapter 6.2 (System Earthing).

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

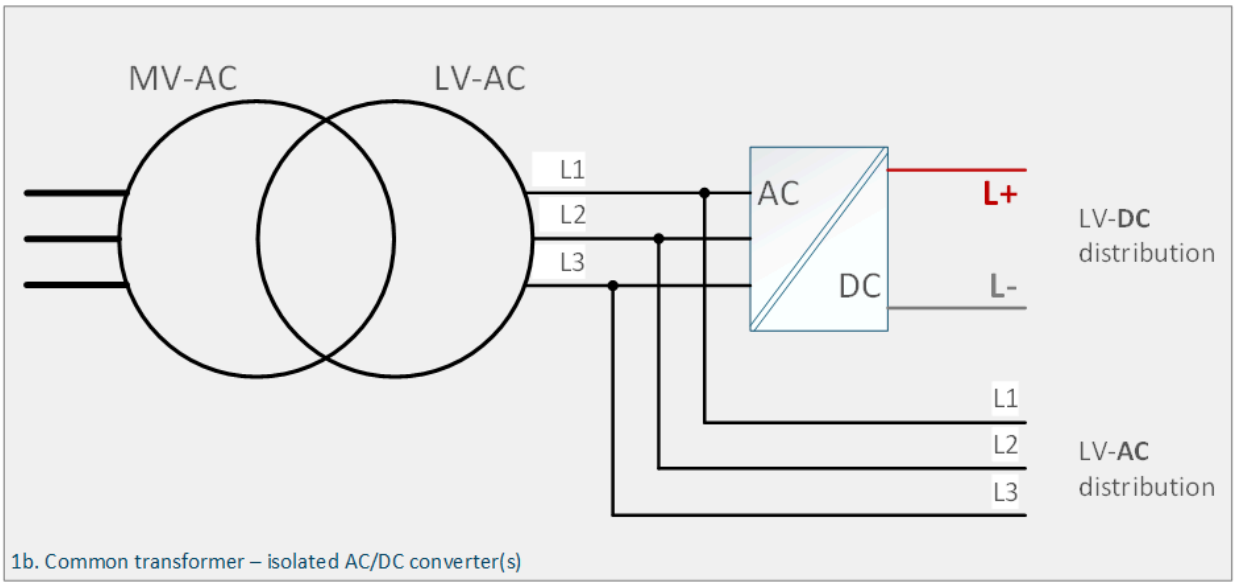


Figure 4.2. Isolation at the central AC/DC converter

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

- Option 1c: shown in Figure 4.3 Separate transformers for DC and AC distribution with separate transformers for the DC distribution and AC distribution, respectively.The AC/DC conversion can be isolating or non-isolating (shown in non-isolating) since the isolation to MV is granted by the MV transformer.

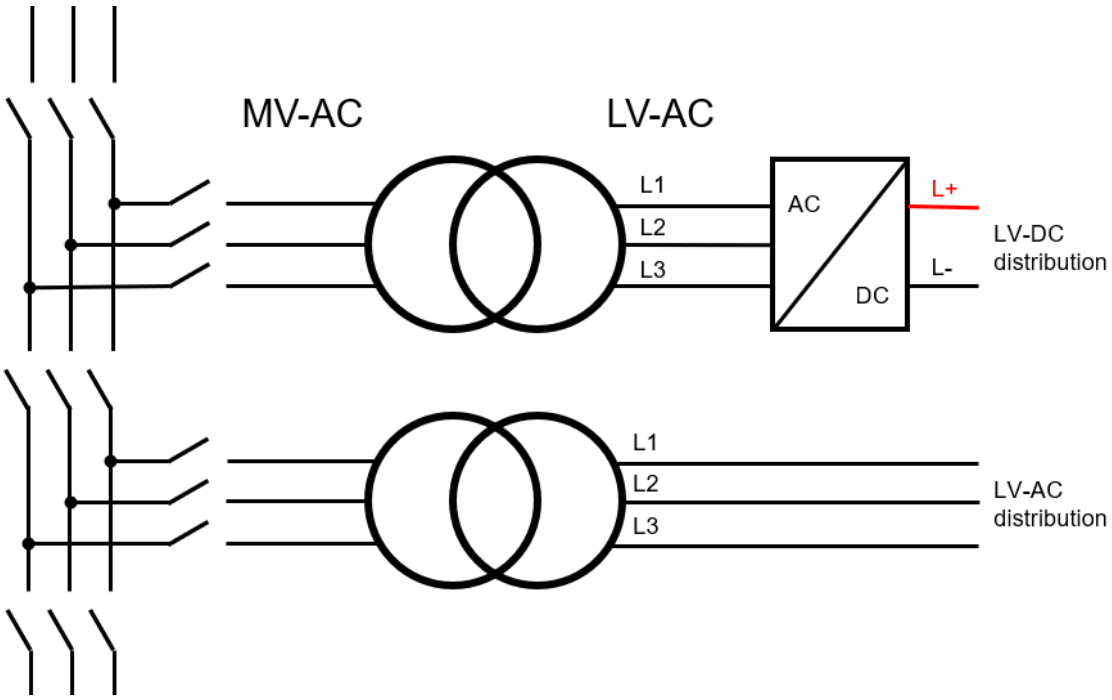


Figure 4.3. Separate transformers for DC and AC distribution

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

- Option 1d: shown in Figure 4.4 with SST for the LV DC distribution part:

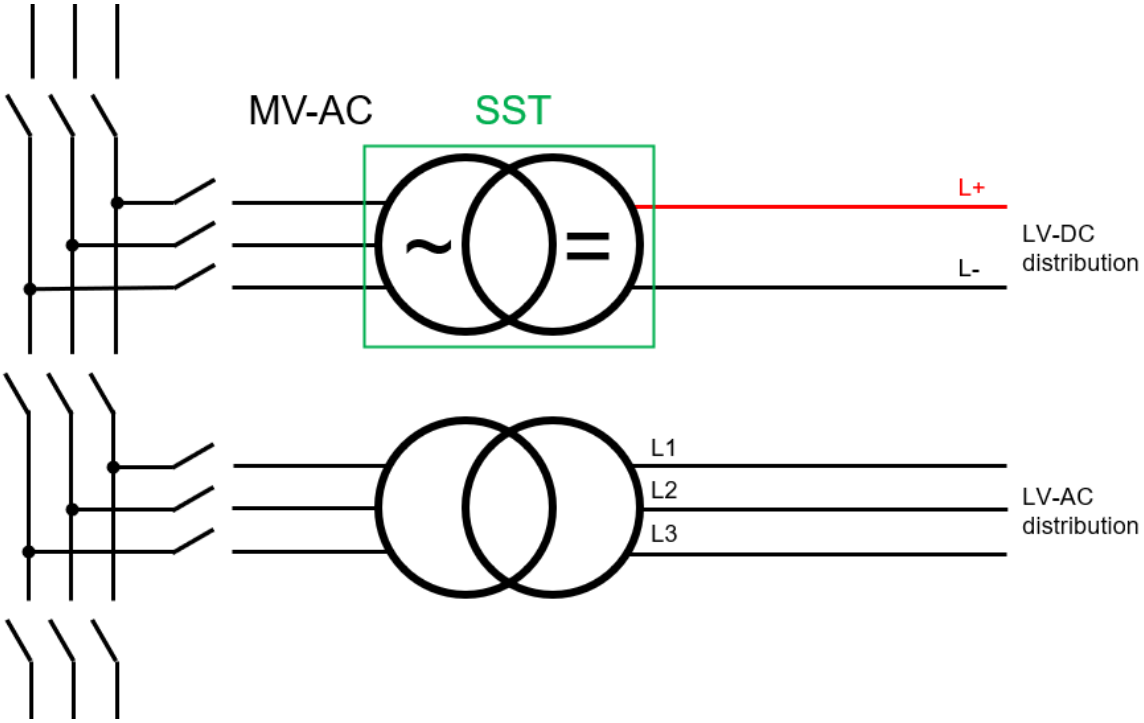


Figure 4.4. Solid-state transformer for DC distribution and transformer for AC distribution

For the SST solution, the SST can either be located remotely from the IT rack with LV DC distribution, or placed closer to the IT rack via MV distribution.

Note that the MV distribution diagram is not shown here, and SST or AC LV transformers may be connected in a different loop depending on the design architecture.

SSTs may have either a single output port or multiple output ports. Figure below shows the example of the conceptual diagram with LVAC distribution, LVDC distribution, and MVAC distribution.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

4.2.2. Stage 2

Full Block 3 IT module (cooling distribution units CDU, Network, IT racks) powered by direct current is shown in Figure 4.5 DC distribution for IT rack and the associated cooling devices, the building infrastructure still runs on AC.

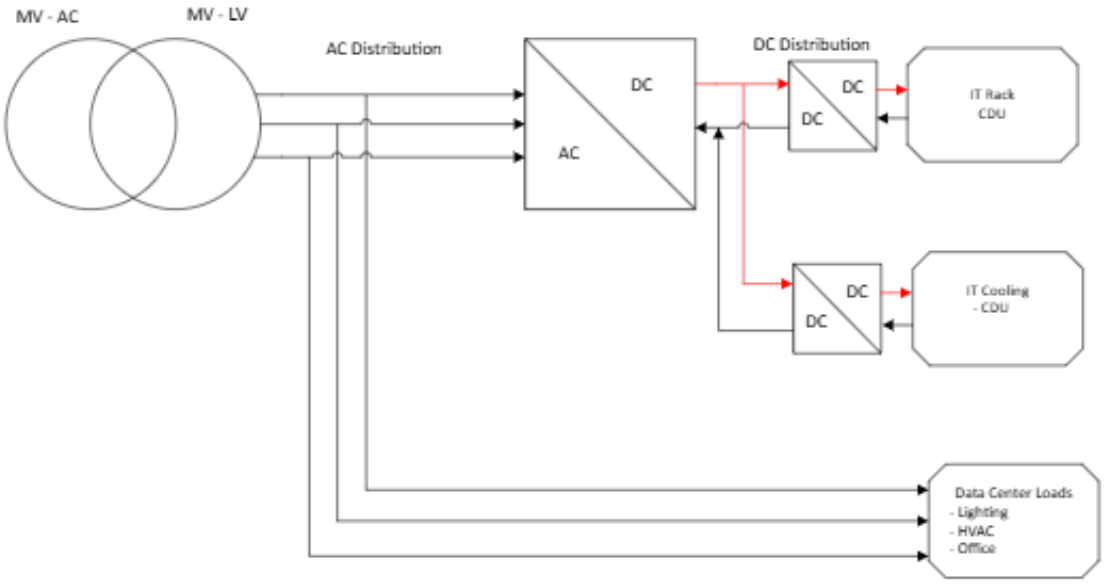


Figure 4.5. DC distribution for IT rack and the associated cooling devices

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

4.2.3. Stage 3

Full Direct current distribution between MV and all equipment in the Data Center. All Equipment is powered by DC only, as shown in Figure 4.6.

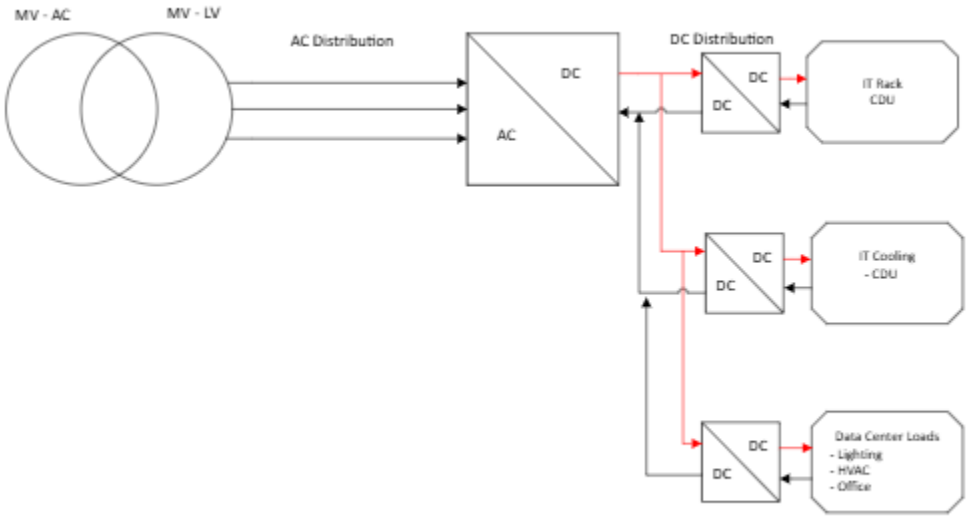


Figure 4.6. Full LVDC distribution for the datacenter, fed by a transformer

Instead of the transformer, it is also possible to install an SST (typically isolated). Please refer to the main components chapter for more details.

At this stage, all loads in the facility will be fed from and connected to the LVDC distribution lines. Supplying facility loads, such as heating, cooling, ventilation, and lighting, has been implemented and proven in reality across many examples, see [9].

Facility loads such as heat pumps, motors for ventilation, use standard components, typically AC asynchronous motors, that are operated with frequency inverters. In a direct current distribution system, the first stage of the inverter, line-frequency AC to the DC link is not needed. Rather, the inverter consists only of the DC/AC conversion step.

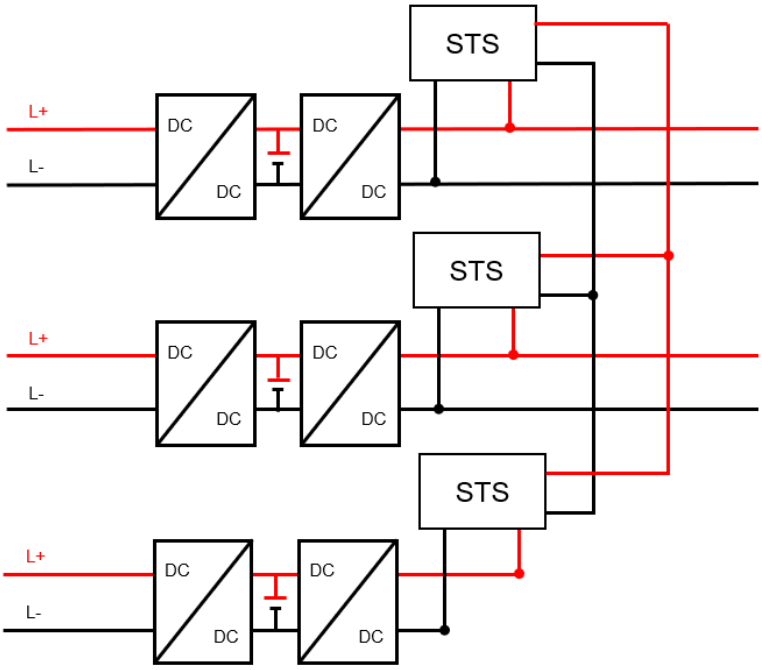
4.3. Redundancy Levels

Currently, the targeted reliability level is 99.9% for training AI data centers and 99.999% for inference AI data centers. Generally, power electronics are the components that have the

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

highest failure rate in a data center (e.g., power shelves, UPS) and this has been managed by a few strategies such as N+1 power blocks, and maintenance bypass.

However, DC distribution does not allow for a maintenance bypass strategy of an uninterruptible power supply, so there needs to be redundancy at the power block level and potentially at the power distribution level as well.



Reserve bus topology (N+1) for DC systems where the DC voltage on both sides of the converter is different. The redundant feeder is the one on the bottom.

If not, there needs to be some kind of upstream DC Static Transfer Switch.

Even for 99.9% availability level, there needs to be some kind of catcher system (type of uninterruptible power supply) to prevent single point of failure at the AC/DC conversion block.

To achieve 99.999%, if the AC/DC conversion block has a high failure rate, then there needs to be redundancy at the AC/DC conversion block level. If there's modularity within the AC/DC conversion block, redundancy at that level may reduce the redundancy requirement at the AC/DC conversion block level.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

Achieving higher levels of availability and concurrent maintainability requires some redundancy at the distribution level. If concurrent maintainability is not required, the system may still meet the target availability but could retain a single point of failure.

Several redundancy schemes can be applied, including block-redundant (catcher architectures) or distributed redundant systems. Each approach has its own advantages and tradeoffs.

Other points of consideration:

- Failure rate dependence on load and ambient temperature
- Infant mortality, and wearout failure rates (bathtub curve): Useful life of the data center is ~10 years - so need to understand how the failure rate changes over that period of time; including on-site spare equipment for repairability
- MTBF and MTTR values, that may depend on the repair strategies like 24hr repair strategy, spare parts provisions or complete system swap.
- Maintenance strategies that can be preventative maintenance, predictive maintenance, condition based maintenance or corrective/emergency maintenance.
- Enhanced monitoring of equipment to predict failures and abnormal wear can increase reliability exponentially.
- For applicable equipment, need bypass test to detect fault that caused the failure of the power module
- The migration to solid-state solutions may require a new, strategic approach to comprehensive coordination studies. Include higher levels of hierarchy as this allows for more flexibility in design.

As data centers migrate towards LVDC distribution, there's still the expectation that data centers continue to maintain similar levels of availability. LVDC distribution introduces new technologies and architectures that require new ways to achieve this level of availability. Another example is that major building blocks (SST, UPS, BBU, PCS) are migrating towards solutions made up of even lower-level modular sub-blocks which must have hot-swap or semi-hot-swap capabilities (Example) A more in-depth look at strategies to achieve availability in LVDC data centers requires further investigation.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

4.4. Power distribution redundancy schemes for DC data centers

In data center engineering, ensuring continuous uptime is a balancing act between cost, complexity, and risk. To evaluate how a data center protects its workloads, we can use a black box approach.

Instead of getting bogged down in the internal mechanics of a specific lithium-ion battery, a complex chiller compressor, or a generator's fuel injector, we treat every piece of infrastructure as a simple, self-contained unit (a "black box") with defined inputs, outputs, and failure rates.



In the figure above, the black box supplies power to the racks. In case the reliability of the box is at par with the required reliability and uptime, this system will comply. This means for a 99.9999% reliability that the black box practically never fails (MTBF = 10E6 years).

However, this is not realistic and therefore redundancy schemes are used, for example below a 2N redundancy topology.



In this case the failure rate becomes the multiplication of the failure rate of each black box. To achieve 99.9999% of reliability, each black box can have a failure rate of 0.1%/y. In an unrepairable scenario, this means both boxes have an MTBF of 1000 years, and combined the expected time for both boxes to be out is 1500 years. In the repairable scenario, where the time to repair is e.g. 24 hours, the probability of having both units down at the same time is 0.548E-8, giving the system an MTBF of 182 million years (for a complete breakdown of both black

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

boxes). Would the boxes only be working in a reserve redundancy scenario, the reliability of the switch and the switching time need to be incorporated in the total reliability of the system. This can increase or decrease the reliability and uptime.

To increase the reliability of a black box, it can contain in itself redundant systems. E.g. a power node may have duplicated the amount of transformers (2N topology) to improve the overall reliability of the power node.

In what follows we will show the redundancy topologies of the MV power distribution to white space systems with a merely comparative indication of reliability and uptime values. However, redundancy topologies may be applied at various locations of the data center, e.g. rack feeds, GPU's, liquid cooling, which when considered a black box, will yield similar results given a failure rate, MTBF, MTTR,…etc. The topologies shown are typical but plenty more different ways of creating redundancy exist.

4.4.1. Background

Most data centers are built around a redundancy scheme on the AC side. Typically used are 2N or N+1 schematics where loops are used to feed the different rack-rows as in figure 4.7.

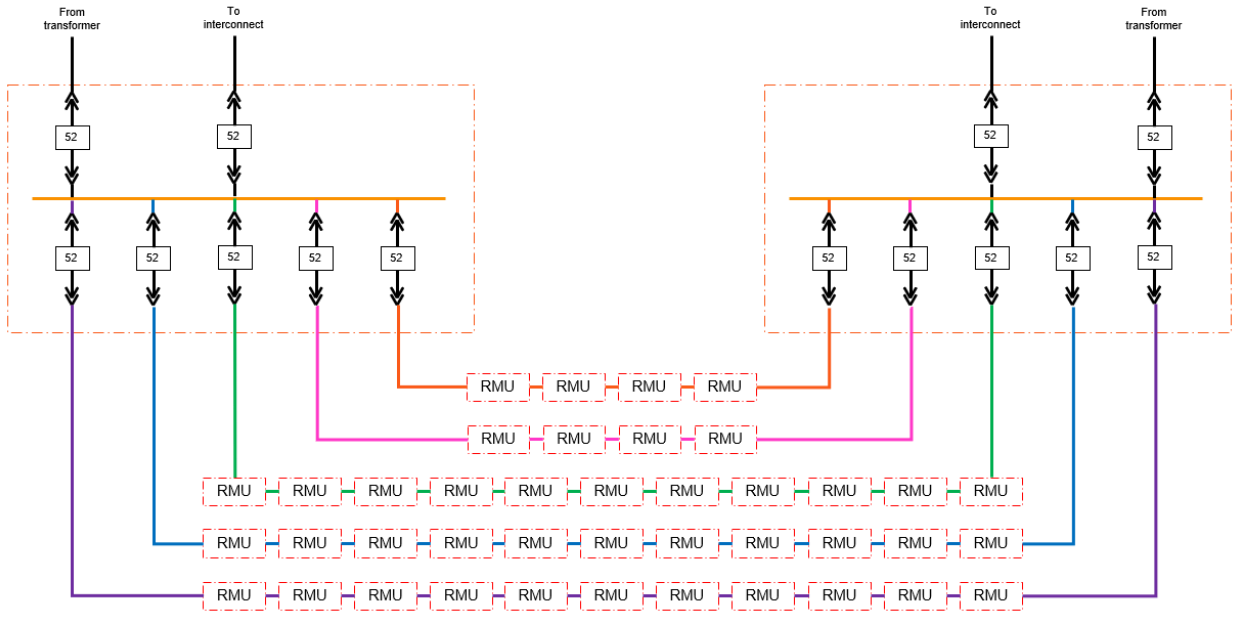


Figure 4.7: example of 2N redundant MV AC power distribution schematic. RMU stands for ring main unit.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

Each ring main unit will be connected to at least one SST or transformer with PCS to provide the LV DC power to the rack row.

Each ring of RMU's can be fed from both sides, allowing to feed the load through the two busses as needed, but with the ring open in between two RMU's as typically it is not allowed to parallel loads by the utility.

In case of maintenance or breakdown of the source or main distribution switchgear, the nodes can be fed through one bus. Each bus is connected to two sources, the main transformer directly connected to the bus or a transformer that can be fed through an interconnecting bus.

An RMU with a transformer(s) or SST(s) is sometimes called a power node, power center or alike, and a typical configuration is as per figure 4.8. In this paper we will call it nodes going forward.

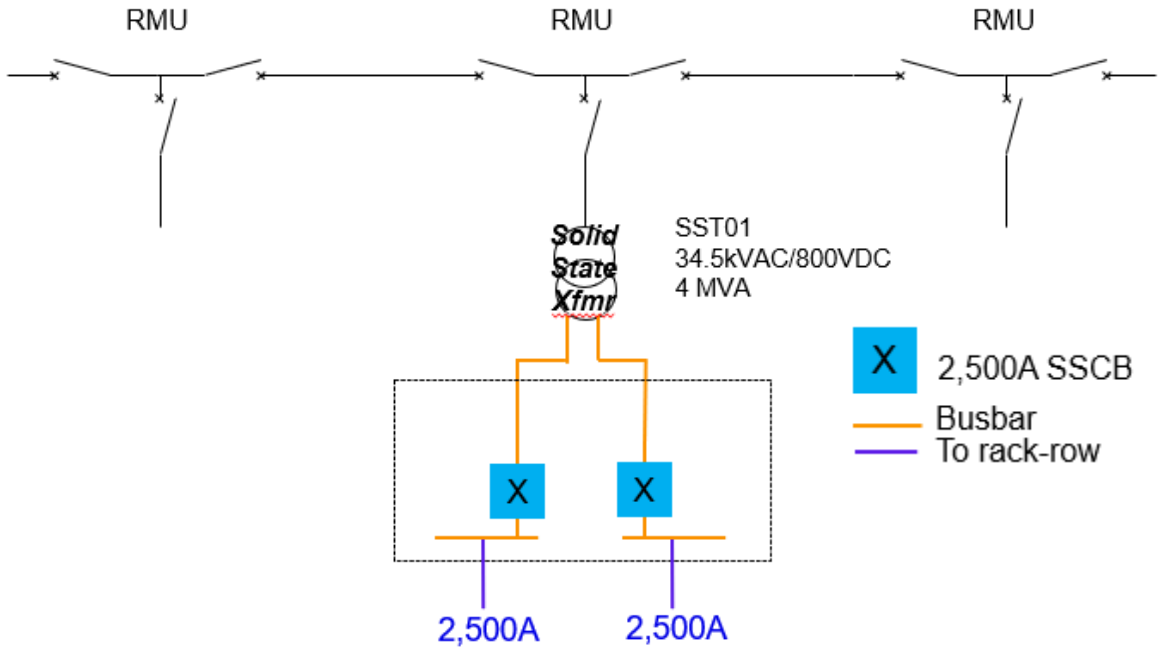


Figure 4.8: schematic showing typical RMU configuration which consists of two line switches to create the loop and one or more T-offs that feed the transformer or SST that will feed the rack-rows, in this case 2 with 2,500A each.

LV switchgear may be included with the power node however it is often placed inside the data center hall, while the MV equipment is often placed outside.

Paralleling nodes is not recommended as this will increase the short circuit currents in the system in case of a fault. This description will not consider schemes with paralleled nodes.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

4.4.2. Reserve or catcher redundancy topology (N+x)

In a reserve redundancy topology (N+x) the needed quantity of power nodes in each ring is complemented with x amount of power nodes to achieve the 99.9999% reliability. For a power node annual failure rate of 0.6% and a MTTR of 6 months, this typically relates to 2 extra nodes to be needed as in the example in figure 4.9.

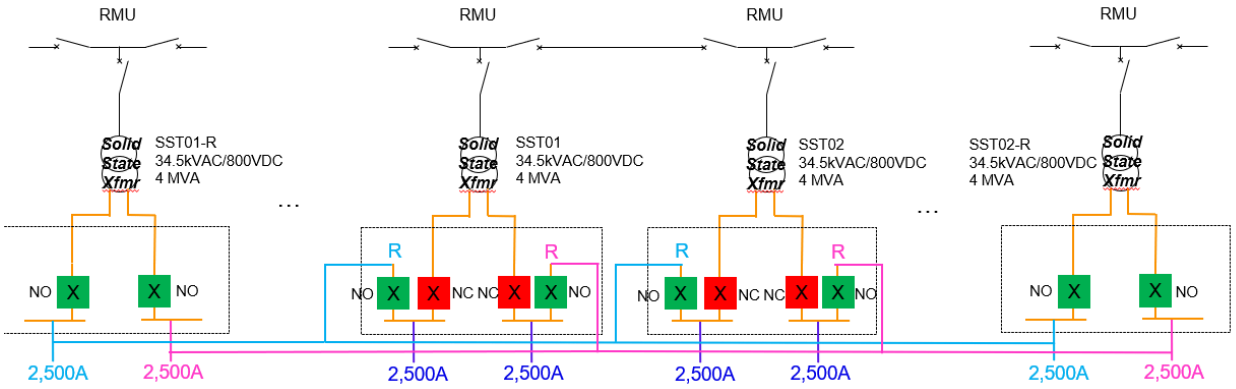


Figure 4.9a: reserve bus configuration with 2 extra nodes (placed for clarity on the sides). One reserve bus set is in pink and the other is in blue. Red switches represent switches in the closed position, green in the open condition.

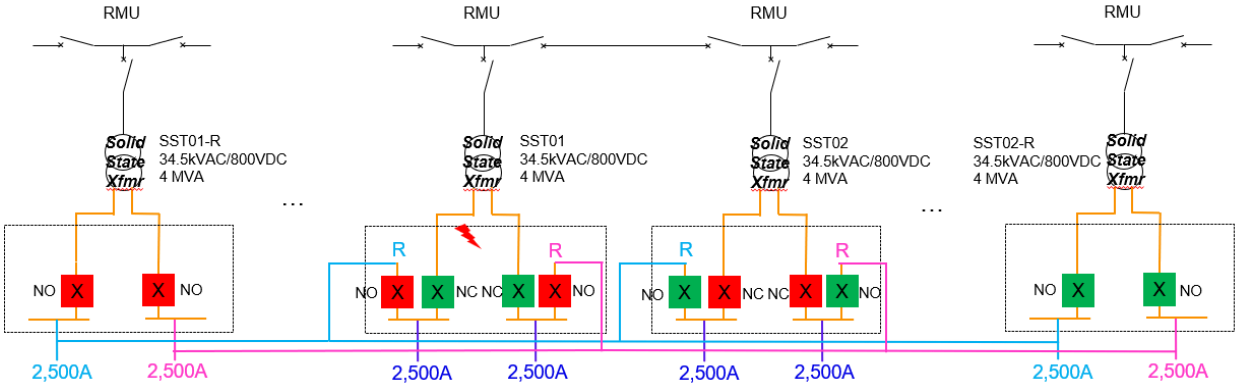


Figure 4.9b: node SST01 being out, shows how node SST01-R picks up its load.

The topology consists in a redundant bus that runs across all the racks and connects to the rack-rows through an additional switch (SSCB) in the node. This topology allows for an automated switching procedure to reestablish power in case of a failure or planned outage upstream of the LV switch. Under this topology it is possible to repair or replace the RMU, SST and in some cases the main LV DC switch while ensuring continuation of power to the rack row.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

The disadvantage of this topology is that it might require a high number of redundant nodes to achieve the requested system reliability if the node failure rate or time to repair is higher. It also does not provide redundancy for the busbar to the rack-row.

4.4.3. Dynamic bus topology (1.5N or 2N)

In a dynamic bus topology (1.5N or 2N) each power node has the ability to cover 1.5 to 2 times the needed capacity of the rack-row, to cover for failed units in the ring. 1.5N means that two nodes will cover for one being out, 2N means one node will cover for one being out. The reliability of this setup can be extremely high, but can also be lower than an N+x strategy. It all depends on how the nodes back each other up.

A first possibility is outlined in figure 4.10. Here the adjacent power nodes can support a node that is out. This schematic has the same amount of switches as the reserve bus topology of figure 4.9, but without the additional nodes. The SST needs to be able to handle the additional power, which means that it needs to be rated 1.5 times to nominal node, to cover for the additional power.

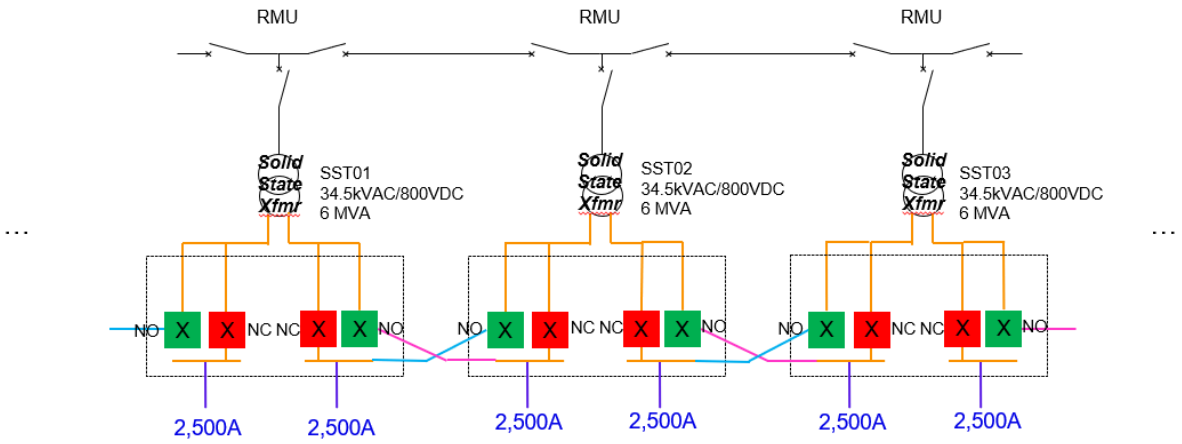


Figure 4.10a: each node is able to support one of its neighbor's rack-row in case of failure. Two nodes on each side can replace one node's outage and nodes need to be designed for 1.5 x power.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

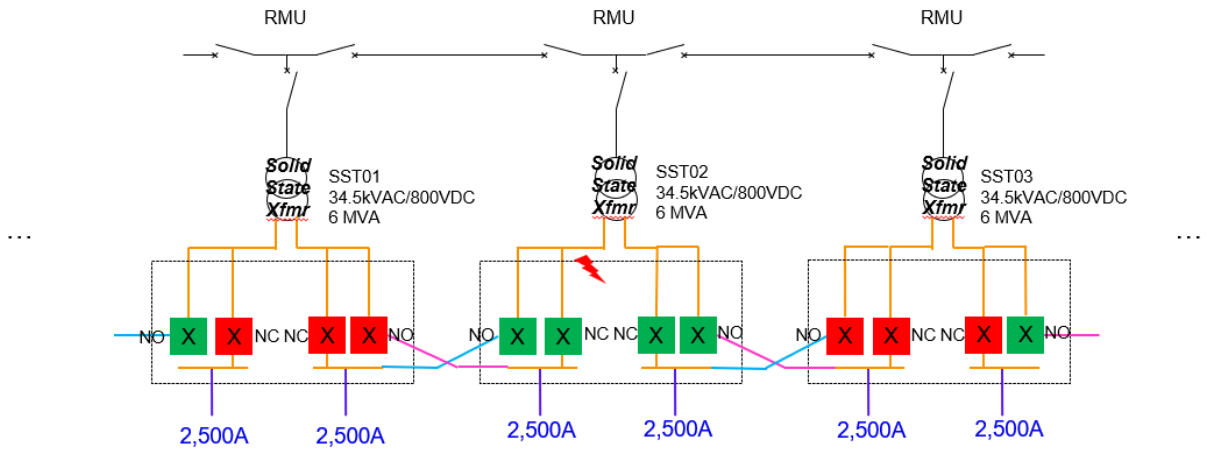


Figure 4.10b: middle node is out and the neighboring nodes pick up each half the load

The scheme will work fine as long as no two nodes fail that are located next to each other. For this reason, this scheme will have a lower reliability than an N+x scheme that does not have that constraint. For a node failure rate of 0.6%, a reliability of 99.9999% cannot be obtained. In fact, the reliability will keep dropping as more nodes are in the loop.

In order to remove the restriction that two neighboring transformers failing would cause an issue, the topology needs to allow backing up a non-neighboring unit, which is achieved in figure 4.11 with redundancy busbars.

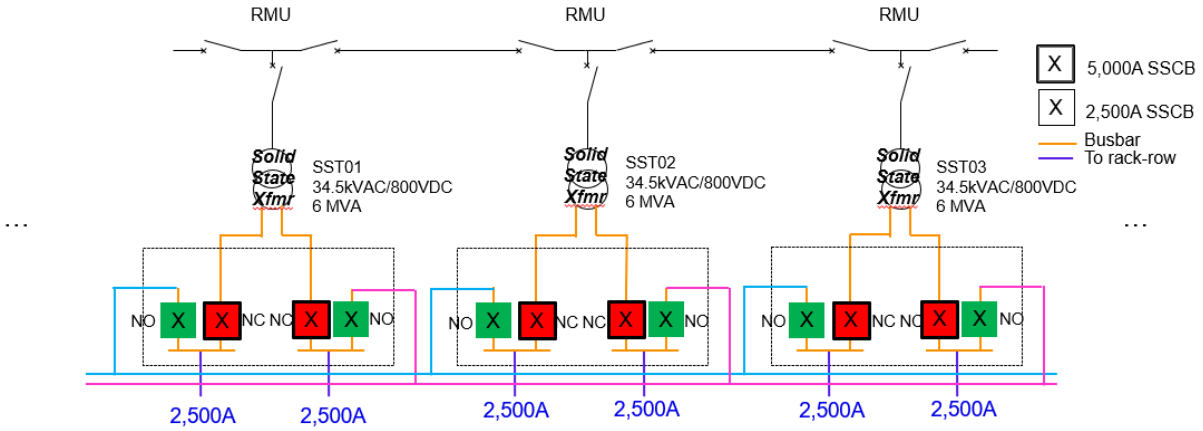


Figure 4.11a: adjusted architecture that allows non-neighboring nodes to support a single node outage.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

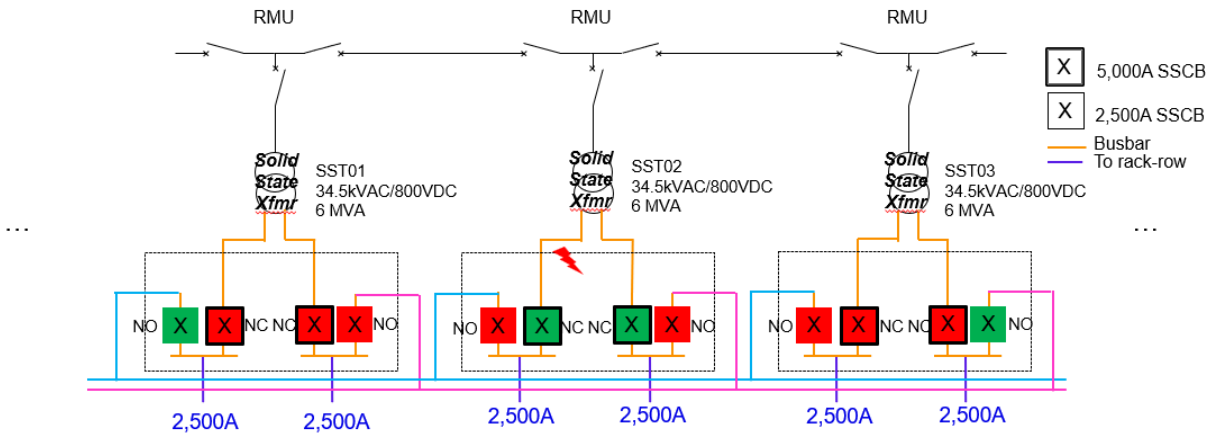


Figure 4.11b: switch positions for an outage in the middle node. A proper procedure is needed to avoid overloading the redundancy busbar in pink and blue. They may be sized larger than the rack-row bus to accommodate for this.

Each node can feed a common bus. As can be seen in figure 4.11, this will require the main switch and associated busbar to be capable of providing twice the amperage, and interlocking will be needed to avoid overloading the SSTs, i.e. only one 2,500 A switch may be open at any time in the node. This restriction is not applicable if the SST is dimensions for twice the power, which would make this a 2N setup.

In case it is not desirable to size the main switch for twice the rack-row power, a topology with additional switches can be used as in figure 4.12.

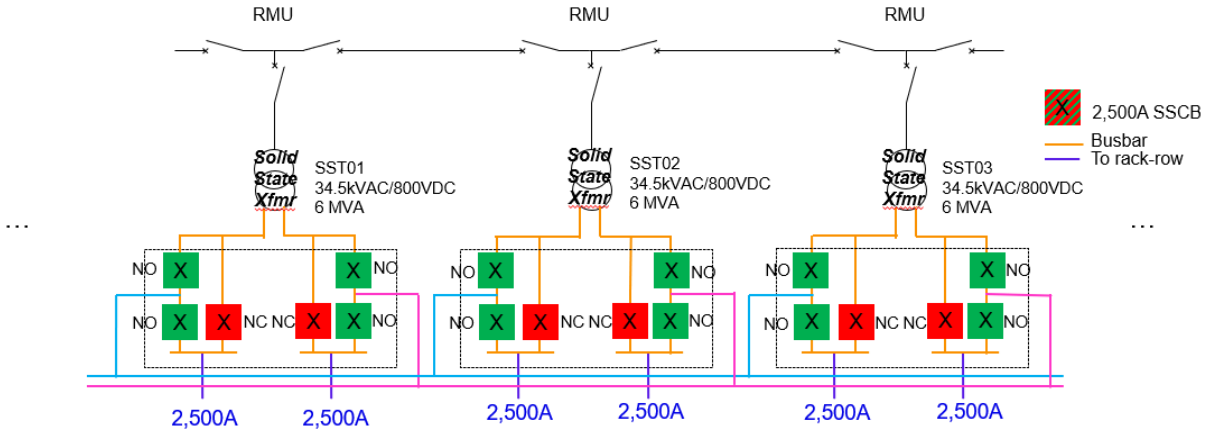


Figure 4.12a: redundancy similar to fig 4.11b but with switches of all similar capacity.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

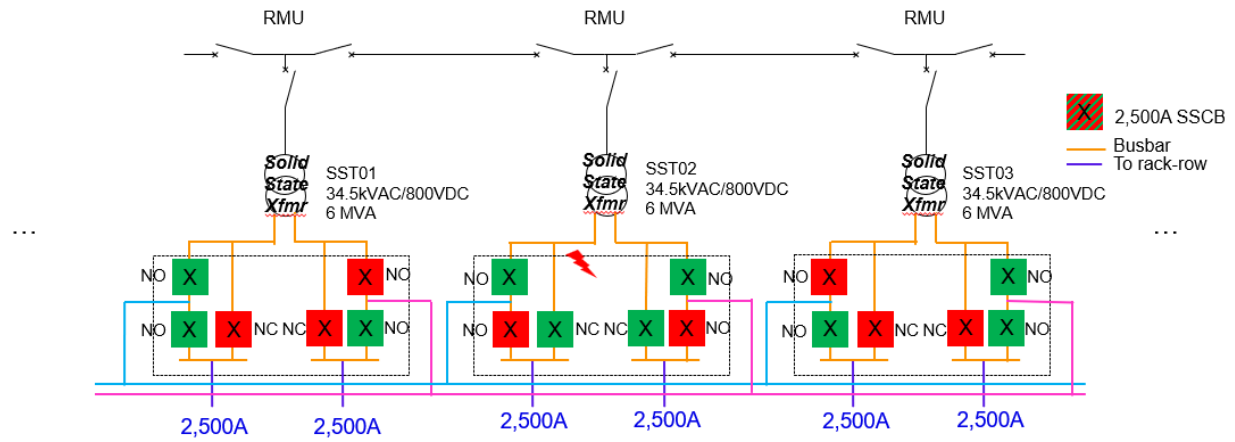


Figure 4.12b: switch positions for an outage in the middle node. The top switch will close at the supporting nodes and the bottom switch will close at the receiving busbar. At no time will both be closed.

Figure 4.12 requires additional switches, but this may have the advantage that none of them has to be sized for twice the capacity and have eventually different parameters for protection when covering for another node. The target reliability of 99.9999% is achieved for an annual failure rate of 0.6% when N ≥ 4. So min 6 units are needed in case of 1.5N, or min 4 units in case of 2N.

In some data center layouts, it is possible that 2 different rack rows are in line with each other and fed by nodes on opposite sides. It is possible to use this layout for redundancy limiting the amount of switches needed, as shown in figure 4.13. When one node is out, the feeder switch is opened and the interconnecting switch is closed. One node's load has doubled.

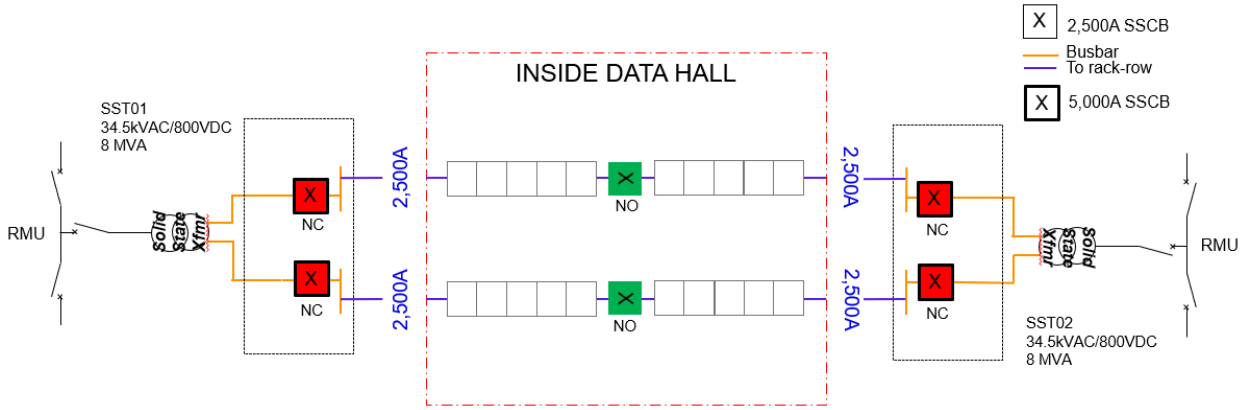


Figure 4.13a: data hall with racks fed from opposite sides. The interconnection is open under normal conditions.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

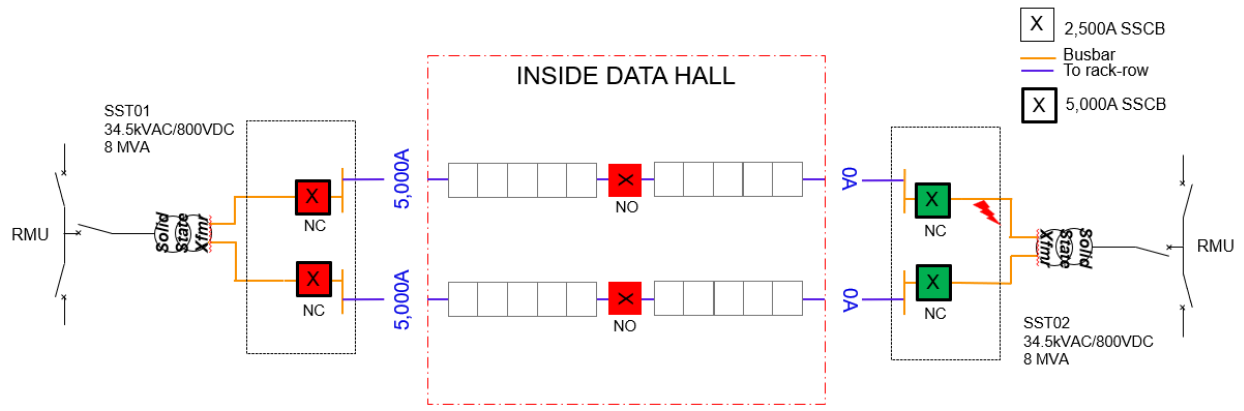


Figure 4.13b: Outage of one node will be covered by the opposing node by closing the tie switch. Busbar and node switches to be rated for twice the amperage.

As can be seen in figure 4.13, the switch at the node is double in size as the interconnecting switch. The SST is also double in size. For this schematic not to fail, both transformers on opposite sides of the rack-row should not be out at the same time. Like in the schematic of figure 4.10, the reliability goes down with a higher amount of nodes in the loop.

In order to avoid the double sized switches, the schematic can be changed into the layout of figure 4.14. This will also require an increase in the amount of switches. Each busbar is connected to a separate switch at the opposite node. When a node is out, the feed switches to the other node. The SST still needs to be sized for twice the power, but the switches do not.

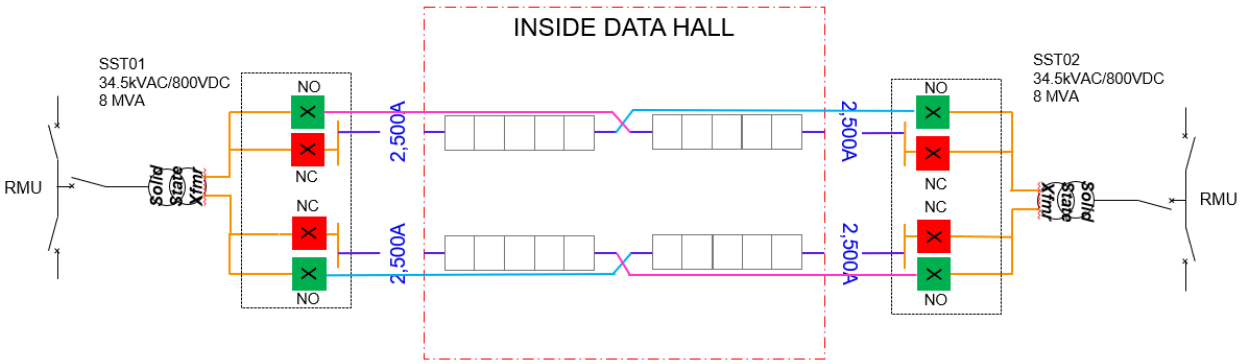


Figure 4.14a: redundancy schematic similar to figure 4.13 but with additional circuit breakers and busbar to avoid higher amperage equipment.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

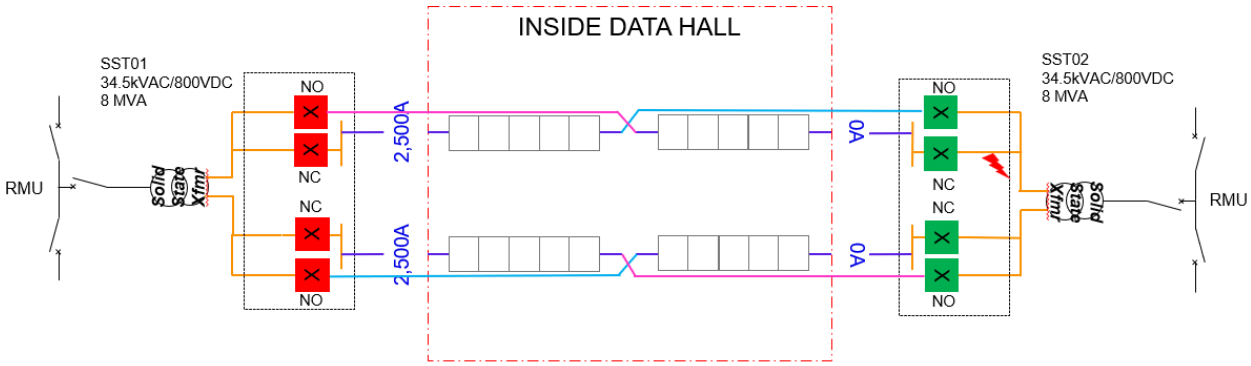


Figure 4.14b: switch positions during outage of one of the nodes.

For all the schematics shown from figure 4.7 to 4.14 the feeder busbar to the rack-row was not redundant: a failure on the feeder busbar cannot be resolved by switching over the source. This can be resolved by adding additional feeder busbars at the node and separating the circuits as shown below in figure 4.15.

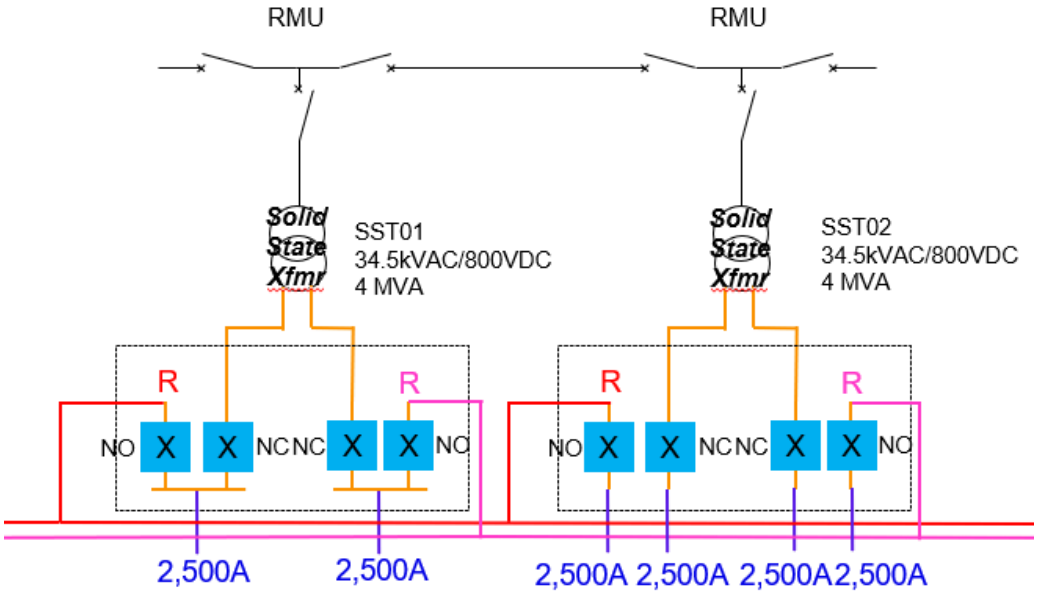


Figure 4.15: example of decoupling redundant circuits from main in power node: on the left the topologies where the feeder busbar is not redundant, on the right with redundant feeder busbar; this modification can be applied to all schematics discussed earlier.

When a failure occurs in the feeder busbar, the redundant circuit takes over immediately.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

4.4.4. Dynamic redundancy topologies at rack level

In the previous discussions the redundancy was achieved by switching power sources on and off in order to achieve continuous operation of the load, whether that is IT or mechanical. Fast switching is definitely required when UPS systems are not available anymore, and to reduce inrush currents, all buses will need to be at voltage at all times. A bus without voltage will need to load up capacitive energy before being able to supply power.

In the previous chapter, N+x reserve redundancy was discussed, where the additional x units are only engaged in delivering power after one or more units have failed. We will discuss this here the use of dynamic redundancy which are typical for rack level redundancy for the example of a 4 to make 3 setup.

In dynamic redundancy, all N+x power sources are on at all times and share the N loads. With the loss of x power sources, the N loads still will be able to perform their tasks. Figure 4.16 shows a possible setup for 4 to make 3.

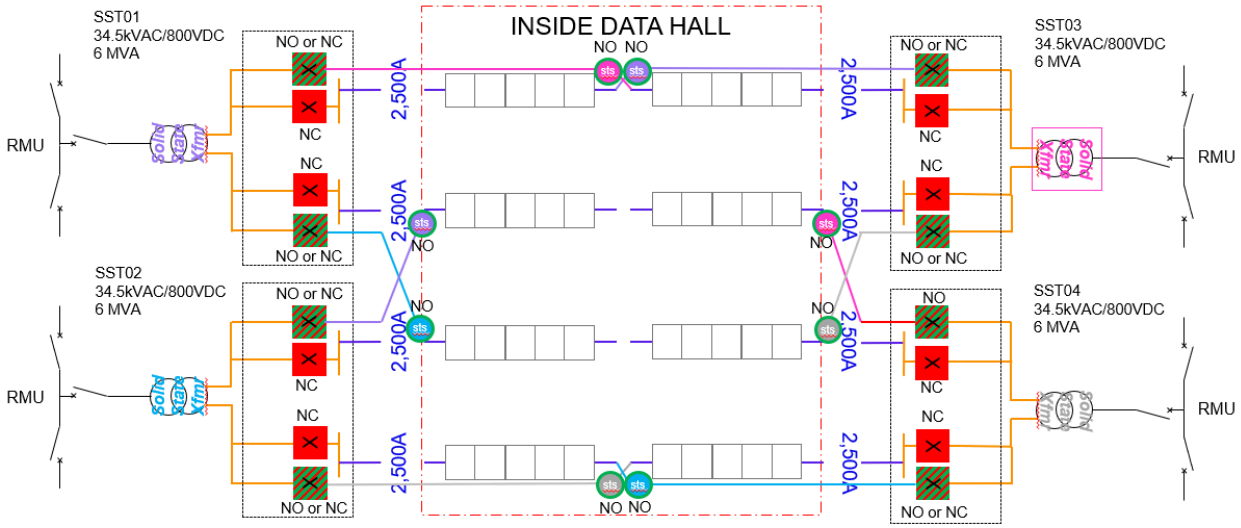


Figure 4.16a: 4 to make 3 setup for power supply to a data center. Note the optional introduction of solid state transfer switches to allow for a fast connection of de-energized busbar. All nodes can also all be on the same side of the data hall.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

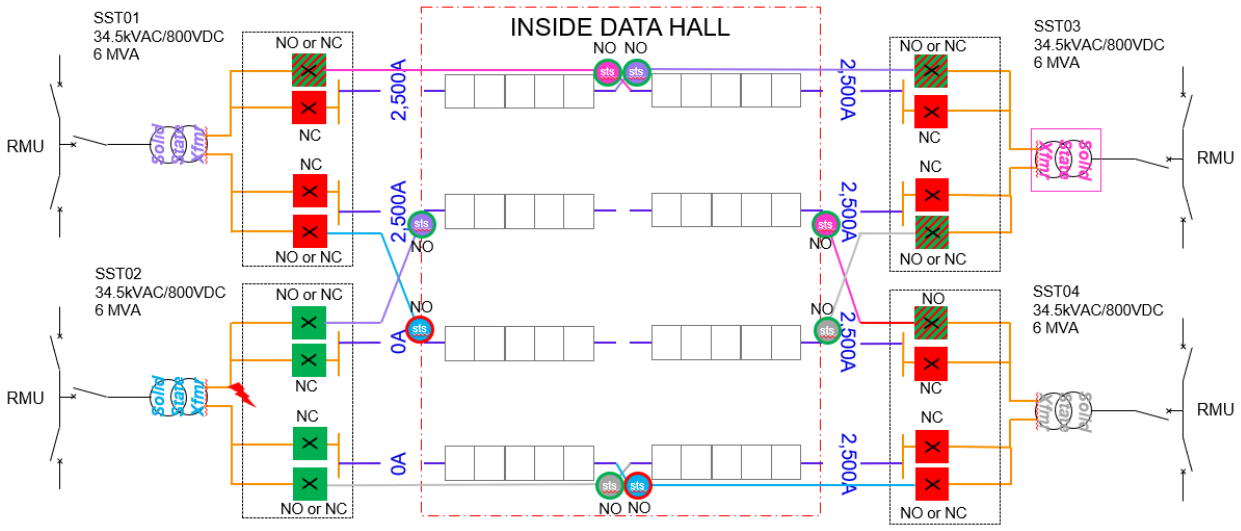


Figure 4.16b: switch positions in case of blue nodes being out.

In figure 4.16, the redundant lines that cover for a certain outage are color coded in the same color as the node they would replace. E.g. if the blue SST is out, the blue line coming from the gray SST and the blue line coming from the purple SST will be activated, as shown in figure 4.17b. In this example each supporting node will need to be able to handle 1.5 times the power, and the load is not evenly distributed when 3 nodes are operational.

To distribute the load over all three nodes, a dual busbar is typical, that allows switching the individual racks to one of the nodes as per figure 4.17.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

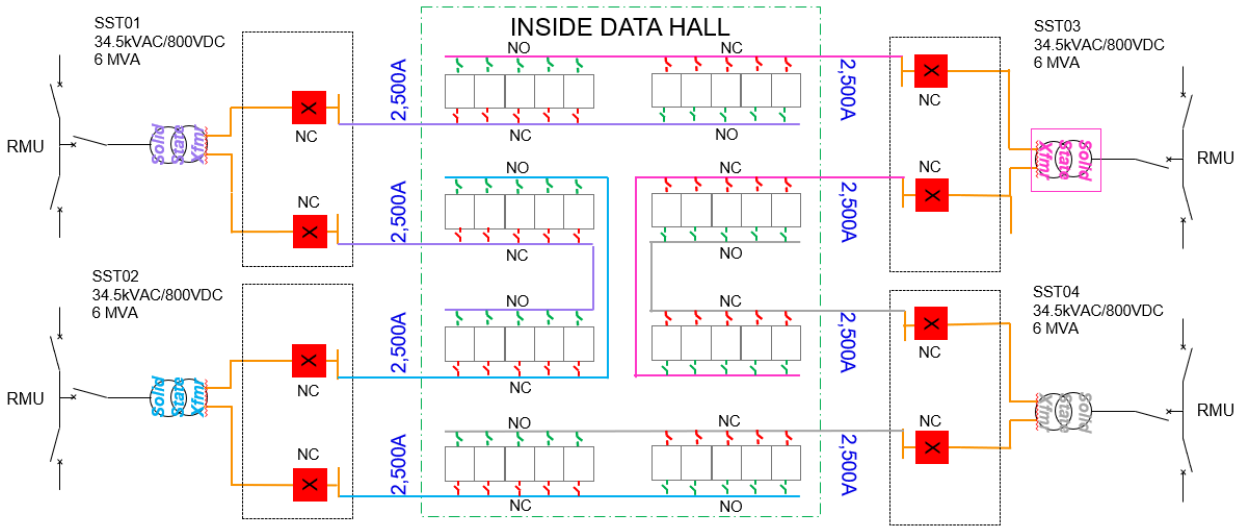


Fig 4.17a: 4 to make 3 with better power distribution in case of a node outage. Note that each rack can be switched to either of the busbars and therefore the amount of switches in the nodes is reduced. All nodes can be on the same side of the data hall.

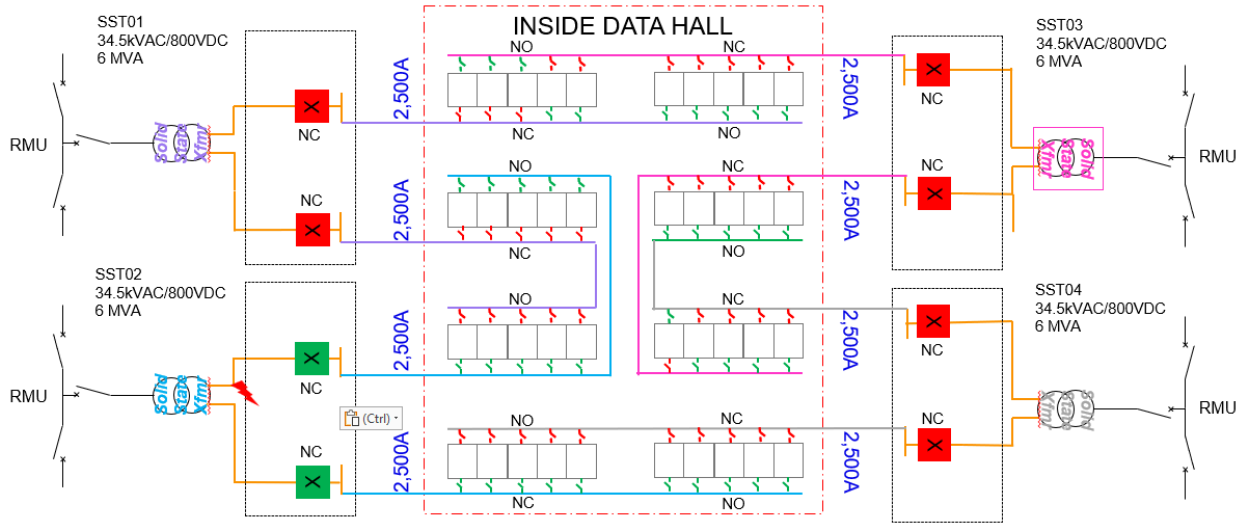


Figure 4.17b: positions of the switches in case of an outage at the blue node.

Figure 4.17 shows the double busbar structure with connection options for each rack to either side of the busbar. In this case after the outage of the blue node, the gray node will supply 14 racks and the remaining two will supply 13 racks.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

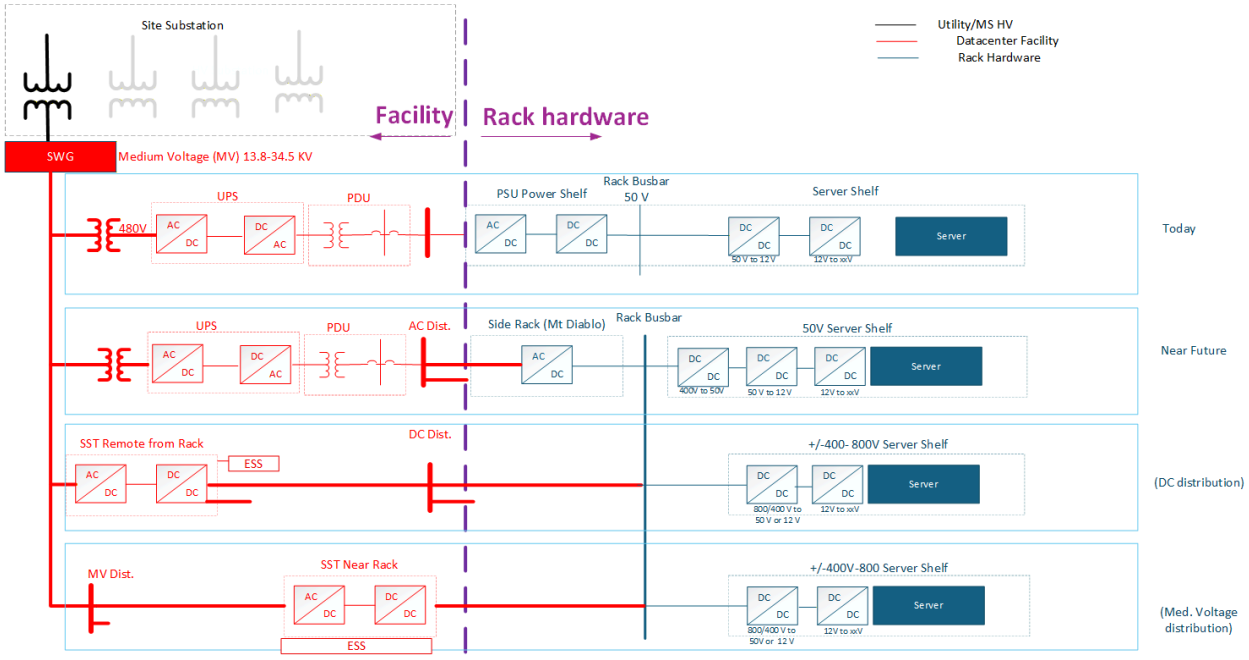


Figure 4.18: LVDC SST vs AC Power Distribution Integration Concept

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

4.4.5. Summary

The table below summarizes the required maximum failure rate of the node and attached busbar for different redundancy schematics, and MTTR budget in order to achieve a 99.9999% reliability of the topology.

System Scheme (N+x to make N)	Redundant Units[a] a x/ (%)	Max Allowable Yearly Unit Failure Rate	Max Allowable MTTR Budget	Tolerance Level / Notes
2 to make 1 (2M1)	1/ (50%)	0.1000%	14.6 Days	High stress on components; tolerates 1 failure.
3 to make 2 (3M2)	1/ (33%)	0.0577%	35.1 Days	Balanced 1-unit cushion configuration
4 to make 3 (4M3)	1/ (25%)	0.0408%	49.8 Days	Strictest component requirements for a 1-unit cushion.
4 to make 2 (4M2)	2/ (50%)	1.6420%	385.8 Days	Most relaxed threshold; allows over a year for repair.
5 to make 3 (5M3)	2/ (40%)	0.9950%	283.6 Days	Highly forgiving 2-unit cushion configuration
6 to make 4 (6M4)	2/ (33%)	0.3694%	225.7 Days	Tolerates. 2 failures; requires under ~7.4 months MTTR
8 to make 6 (8M6)	2/ (25%)	0.2622%	160.1Days	Tolerates 2 failures; requires under ~5.2 months MTTR
10 to make 8 (10M8)	2/ (20%)	0.2035%	124.1 Days	Tolerates 2 failures; requires under ~4.1 months MTTR.

[a] the % indicated in this table gives an indication of the extra cost of the redundant system compared to no redundancy. E.g. in a 6M4 33% of the cost is to build the redundancy

The most common failure rates for switchgear and transformer combinations are 0.6% to 1.2% depending on the quality of the product.

The reliability of a N+x reserve redundancy is slightly higher than that of an equivalent dynamic redundancy as the units that are redundant are not operational and less prone to failure. However, when called on, the loads on the units increase immediately which can cause the node to fail.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

4.5. Inclusion of redundant power generation

Figure 4.19 shows the possible locations of redundant power generation and/or storage to support outage of a node or part of a node. The closer the generation is to the data center load, the less it depends on possible failing items from the power node. E.g. If option 1 is applied, an outage caused by the SST will render the power generation unusable. It is also recommended to have enough capacity to run the power needed by the load at any time and provide enough energy to allow for a safe shutdown if only temporary in nature.

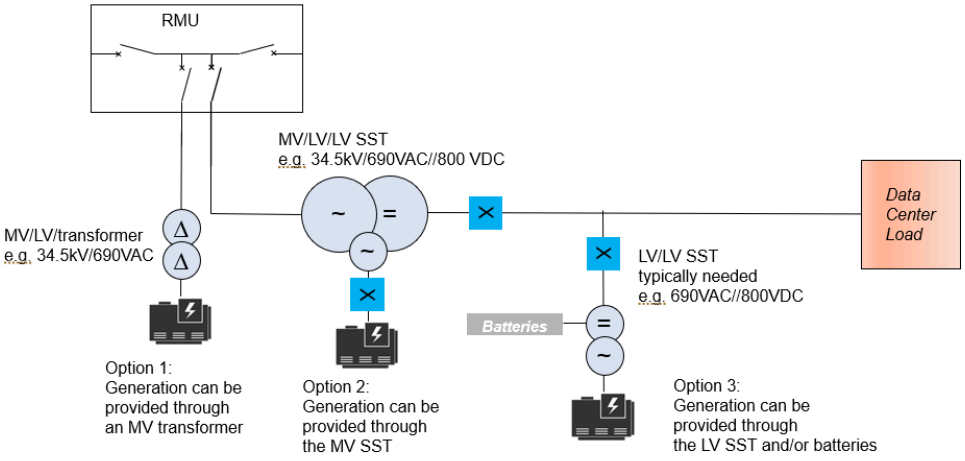


Figure 4.19: generation/storage options for standby power redundancy.

4.6. Verified Benefits of DC installations

LVDC installations have been installed and operating for years. This subsection gives some examples.

- Schaltbau's NExT Factory is a new, completely electrified building in operation since Sep.2023, Schaltbau opens NExT Factory in Velden. It is a symmetrical, 2-wire DC-distribution with normal operating voltage band up to 750 V – installed following the DC-INDUSTRIE description, now available as VDE SPEC 90037. After about 2.5 years of operation, they reported these achievements, Schaltbau's NExT Factory lessons from the world's first DC microgrid powered plant
- DC-powered datacenter. In the year 2012 already, ABB and Green, an IT service provider, opened a datacenter, World's most powerful DC data center online | News

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

center | ABB. They conclude that “DC technology should preferably be used in new and large data centers”. See also pages 16ff of ABB Review 4|2013, ABB Review No. 4 2013

- Production cell of BMW.The car manufacturer has converted a car body cell in their factory in Bavaria to DC supply. After almost two years of operation, 8.8% less energy is needed to perform the same job as previously with AC supply, 'With direct current towards more sustainability in automobile production

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

5. Voltage Levels and Power Quality

This section reviews common definitions, introduces key concepts, and discusses critical considerations around voltage.

5.1. Voltage Classes

While many sub classifications of voltage exist, there are two primary classifications commonly recognized by codes and standards – low voltage and high voltage.

5.2. Low Voltage (LV), Including Sub-Classes (e.g. ELV, LES)

5.2.1. General Definition

In the context of power distribution, low voltage is defined by international product standards and regulations along with installation codes as operation at or below 1000 VAC or 1500 VDC. (The 1.5x multiplier from AC RMS to DC is based on comparable insulation stress considering AC peak as √2 factor above AC RMS voltage.)

- NFPA 70®, the National Electrical Code® (NEC®) - Article 100 Definitions -
  High Voltage: *A potential difference over 1000 volts ac, 1500 volts dc, nominal.*
- IEC 60050-195 International Electrotechnical Vocabulary (IEV) - Part 195 (05-25) -
  Low Voltage (LV): *voltage not exceeding a conventionally adopted limit*
  - *Note 1 to entry: For AC, the conventionally adopted limit is 1 000 V.*
  - *Note 2 to entry: For DC, the conventionally adopted limit is 1 500 V.*
- IEC 60364-1 Low Voltage Electrical Installations - Part 1 - Section 11.2 -
  IEC 60364-1 covers
  a) circuits supplied at nominal voltages up to and including 1 000 V a.c. or 1 500 V d.c.;
- IEC TS 63527 Safe management and operation of electrical installations -
  Low Voltage (LV): *voltage normally not exceeding 1 000 V AC or 1 500 V DC*
  High Voltage (HV): *voltage normally exceeding 1 kV AC or 1,5 kV DC*
- European Union (EU) Low Voltage Directive (LVD) - Article 1 Subject matter and scope:
  "This directive shall apply to electrical equipment designed for use with a voltage rating of between 50 and 1 000 V for alternating current and between 75V and 1 500 V for direct current, other than the equipment and phenomena listed in Annex II."

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

- China GB/T 35727:

  3.6 本标准将3000(±1500)V~±50kV 划定为中压范围,将110V~1500(±750)V 划定为低压范围。

  3.6 This standard sets ±3,000 V (inclusive) to ±50 kV as the high-voltage DC distribution range and 110 V to 1,500 V (inclusive) or ±750 V as the low-voltage DC distribution range.

Even where the NEC® has been updated to permit 2000 VDC for certain solar PV installations (see Figure 5.1), the dividing line between low and high voltage remains at 1500 VDC. This boundary is supported by established insulation coordination practices, equipment design constraints, and appropriately defined personnel protection requirements.

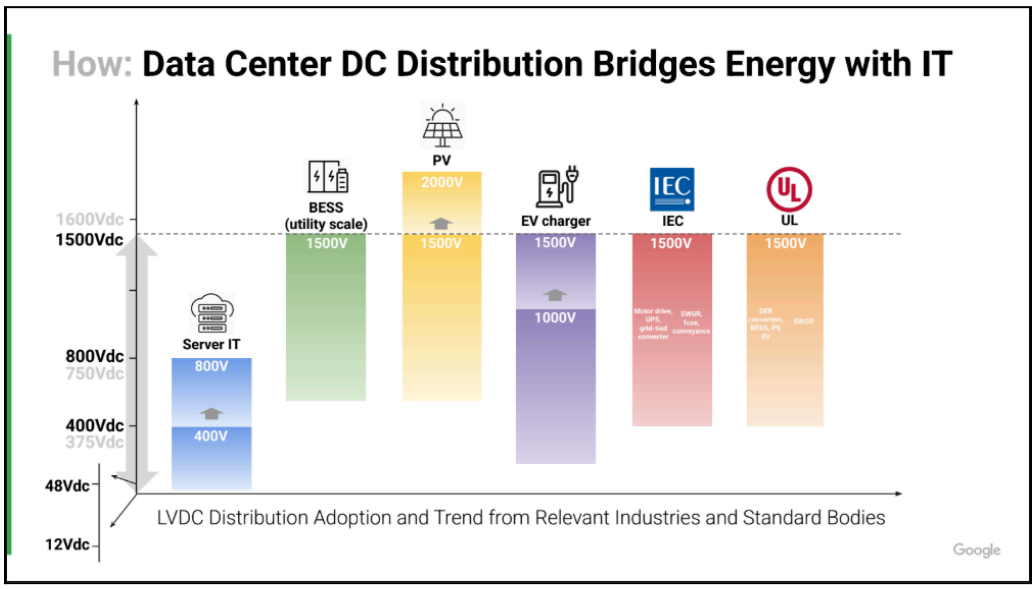


Figure 5.1. Low Voltage DC Limit (Ref.: Intro of Data Center LVDC)

Despite general adherence to this limit, it is worth mentioning that there may be certain local exceptions or product standard scoping limitations which define lower voltages (for more info see Section 11). For example, in Japan, 600 VAC and 750 VDC are commonly defined in JIS standards.

Less commonly, there are also higher definitions applied by certain product standards such as IEEE C37.14 which defines a 3200VDC limit for LVDC power breakers. However, these

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

standards are still bound by local installation codes such as the NEC® which adheres to the 1500 LVDC limit.

From a physics perspective, the possibility of corona discharge (whereby air surrounding energized electrical equipment becomes ionized and conductive) exists for certain conditions of sharp edges, defects, and low pressure environments with voltages as low as 300V. However, for voltages above 1000VAC or 1500VDC, additional design and testing measures become necessary, considering the large air gaps required and decrease in the dielectric strength of air with altitude.

5.2.2. Extra-Low Voltage (ELV)

Within the low voltage classification exists a sub-classification for extra-low voltage defined according to IEC 61140 as ≤50 VAC or ≤120 VDC. Note that whereas LVDC generally applies a 1.5x multiplier between DC and AC rating, ELV applies a 2.4x multiplier in consideration of the electrical shock hazard as defined by IEC 60479-1.

Additionally, basic protection against electric shock may be provided by ensuring touch voltage is limited to ≤25 VAC or ≤60 VDC in dry locations where large-area contact with live parts is not expected (≤6 VAC or ≤15 VDC otherwise) and supplying voltage from a battery, safety isolating transformer, or safety equivalent. The dependence of the exact voltage threshold for ELV on the conditions aligns with the more general definition below. Additional criteria specific to SELV (ungrounded) and PELV (grounded) systems are also identified.

- IEC 60050-195 International Electrotechnical Vocabulary (IEV) - Part 195 (05-24) - Extra-Low Voltage (ELV): *voltage not exceeding the maximum value of the prospective touch voltage that is acceptable to be maintained indefinitely under specified conditions of external influences*
- IEC 60050-195 International Electrotechnical Vocabulary (IEV) - Part 195 (06-28) - SELV (safety extra-low voltage) system: *electric system in which the voltage cannot exceed the value of extra-low voltage:*
    - *under normal conditions and*
    - *under single fault conditions, including earth faults in other electric circuits*

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

- IEC 60050-195 International Electrotechnical Vocabulary (IEV) - Part 195 (06-29) - PELV (protective extra-low voltage) system: *electric system in which the voltage cannot exceed the value of extra-low voltage:*
    - *under normal conditions and*
    - *under single fault conditions, except earth faults in other electric circuits*

Similar though different requirements exist in NFPA 70E, where ≤50 V (AC or DC) is defined in Section 110.2(B) as part of general criteria for an electrically safe working condition. This is extended to ≤100 VDC when working on certain equipment under additional restrictions for fault current (e.g. 100VDC, 40 mA in Sections 320.3, 330.3, 340.4, 350.9). Note that OSHA also defines the 50V (AC or DC) threshold with certain exceptions in 29 CFR 1910.333(a)(1).

5.2.3. Limited Energy Systems (LES)

The NEC® additionally defines wiring classes 1 - 4, with classes 2 - 4 categorized as limited energy for mitigation of shock and fire hazards (see Table 5.1). Details of classes 2 and 3 voltage and power ratings are provided in NEC Chapter 9, Tables 11(A) and 11(B).

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

Table 5.1. Limited Energy System Wiring Classes in the NEC®

	Common Applications	Typical Voltage	Power	Fire Risk	Shock Risk
Class 2	Power over Ethernet (PoE), alarms, automation	≤30 VAC, ≤60 VDC	≤100 VA	Low by design	Low by design
Class 3	medical, audio, industrial	up to 150 V	~150 VA	Mitigated by cabling rules	Low by design
Class 4	Fault Managed Power Systems (FMPS)	up to 450 V peak	100 W (L-E or L-L faults)	Low by reaction	Low by reaction

5.3. High Voltage (HV), Including Sub-Classes (e.g. MV, HV, EHV, UHV)

5.3.1. General Definition

In the context of power distribution, high voltage is defined in contrast to low voltage as greater than 1000 VAC or 1500 VDC. (The 1.5x multiplier from AC RMS to DC is based on comparable insulation stress considering AC peak as √2 factor above AC RMS voltage.)

5.3.2. Medium Voltage (MV)

The immediate voltage range above and adjacent to low voltage is medium voltage. The upper limit of medium voltage is variously defined at different levels, but typically not distinguished between AC and DC.

ANSI C84.1 lists the upper limit of medium voltage as <100 kV. In an informative note for the definition of high voltage, the NEC® identifies that circuits and equipment rated for up to 52 kV are commonly referred to as medium voltage. It specifically defines Type MV cable for ≥2001 V and ≤35 kV as well as other cable types in Table 690.31(A)(3)(1) for ≤2000 V. This corresponds with a typical application of medium voltage in many utility distribution systems.
With continuing advancements in renewable energy generation, additional market demand is bringing additional focus on the development of specific standards for MVDC (particularly up to 2kV and 3kV DC levels) and a recent report has been published by the IEC exploring MVDC grids for an all-electric society.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

5.3.3. High Voltage (HV), Extra High Voltage (EHV), Ultra High Voltage (UHV)

Voltages beyond medium voltage are typically the domain of utility transmission systems. ANSI C84.1 defines:

- High Voltage (HV): ≥100 kV and ≤230 kV
- Extra High Voltage (EHV): >230 kV and <1000 kV
- Ultra High Voltage (UHV): ≥1000 kV

5.4. Nominal Voltage

Nominal voltage identification serves an important purpose to uniquely label an AC or DC system. It also defines both an installation rating class for facility wiring and a product rating class for sources, loads, and prosumers (devices which may be either producers or consumers). Several key definitions are provided as follows (refer to Section 5.6 for further information on voltage bands and levels):

- NFPA 70®, the National Electrical Code® (NEC®) - Article 100 Definitions -
  Nominal Voltage: *A nominal value assigned to a circuit or system for the purpose of conveniently designating its voltage class (e.g., 120/240 volts, 480/277 volts, 600 volts).*
  - *Informational Note No. 1: The actual voltage at which a circuit operates can vary from the nominal within a range that permits satisfactory operation of equipment.*
  - *Informational Note No. 2: See ANSI C84.1-2011, Voltage Ratings for Electric Power Systems and Equipment (60 Hz).*
- IEC 60050-601 International Electrotechnical Vocabulary (IEV) - Part 601 (01-21) -
  Nominal Voltage of a system: *a suitable approximate value of voltage used to designate or identify a system*
- IEC 60050-826 International Electrotechnical Vocabulary (IEV) - Part 826 (11-01) -
  Nominal Voltage <of an electrical installation>: *value of the voltage by which the electrical installation or part of the electrical installation is designated and identified*
- IEC TR 63282 - LVDC systems - Assessment of std voltages & power quality reqs -
  DC system nominal voltage: U_n *, value of the voltage by which the electrical installation or part of the electrical installation is designated and identified*
  - *Note 1 to entry: The DC system nominal voltage U_n is within the nominal band $[U_2; U_3]$ but not always half-way between U_2 and U_3. In all cases $U_2 \leq U_n \leq U_3$*
  - *Note 2 to entry: For a bipolar system, it is recommended to use a dual notation, for example, "$\pm U_{L\text{-}M}$" or "$U_{L\text{-}M} / U_{L\text{-}L}$".*

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

IEC 60038 lists the standard nominal voltage values for AC and DC. At the time of writing it is being amended for LVDC with IEC 60038/AMD2/FRAG1 ED7 CDV (initially approved by national committees and advanced to FDIS stage without changes to nominal voltage), where 175/350V, 350V, 350/700V, 700V, 700/1400V and 1400V are the standard voltages for LVDC from 120V to 1500V. These voltages are consistently scalable with a factor of 2 in this range. Also Current/OS and ODCA (recently) use 700V as the nominal voltage. Therefore 350V, 700V and 1400V are used as the nominal voltage of the DC grids. (400V and 800V would scale to 1600V which is above the low voltage limit of 1500V.)

5.5. Rated Voltage vs. Operating Voltage

The rated voltage of a device (source, load, or prosumer) is visible on the nameplate and used to determine the various voltages applied during qualification testing to applicable standards. It is important to note that this may be more than one value or may be a range of values.

- IEC 60050-151 International Electrotechnical Vocabulary (IEV) - Part 151 (16-49) - Rated voltage range: *voltage range as declared by the manufacturer expressed by its lower and upper rated voltages [SOURCE: IEC 62368-1:2010, 3.3.10.5]*
- IEC 60050-614 International Electrotechnical Vocabulary (IEV) - Part 614 (03-09) - Rated voltage <of equipment>: *voltage value assigned by a manufacturer or other entity for a specified operating condition of a component, device or equipment*
    - *Note 1 to entry: The value for the rated voltage of high voltage equipment is generally assigned from the list of highest voltages for equipment in Tables 3, 4 and 5 of IEC 60038:2009.*
    - *Note 2 to entry: The value for the rated voltage of low voltage equipment is generally assigned from the list of nominal voltages in Tables 1 and 6 of IEC 60038:2009.*
    - *Note 3 to entry: Equipment may have more than one rated voltage value or may have a rated voltage range.*

The operating voltage is closely related to the rated voltage, but differs in important ways. Here below is the official definition of operating voltage according to the IEC:

- IEC 60050-601 International Electrotechnical Vocabulary (IEV) - Part 601 (01-22) - Operating voltage (in a system): *the value of the voltage under normal conditions, at a given instant and a given point of the system*
    - *Note – This value may be expected, estimated or measured.*

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

Importantly for direct current within IEC TR 63282, both a continuous and a temporary operating voltage range are defined with corresponding voltage bands, see Section 5.6 (Voltage Bands). This has the potential to be broadly adopted and is already being incorporated into IEC product standards, especially with the pending update and release of IEC 60038. It is important that the operating rating considers the highest applicable sustained voltage for the product function.

The operating voltage is tightly coupled to the applicable product standard and not necessarily consistent across products. A familiar example is the AC/DC power adapter typically supplied with a laptop computer. Recently IEC 60950-1 has been obsolesced and replaced by IEC 62368-1, but for this example of Table 5.2 and Figure 5.2, a default tolerance for the operating voltage of 90 - 254 VAC applies to the rated 100 - 240 VAC (as no tolerance is indicated on the rating nameplate).

Table 5.2. Example Operating Voltage for Rated Voltage

Rated Voltage	Operating Range (Default Tolerance)	
	IEC 60950-1	IEC 62368-1
	AC: ±10% for 230/400 VAC, -10/+6% otherwise DC: -15/+20%	AC: ±10% DC: -15/+20%
100 - 240 VAC	90 - 254 VAC	90 - 264 VAC



Figure 5.2. Rated Voltage Example (Ref.: Current/OS: DC Distribution for Data Centers)

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

5.5.1. Slash Ratings, Neutral, and Midpoint

For AC, it is common to find 3-phase or 1-phase products with a slash voltage rating where the larger number represents the line-line voltage and the smaller number represents the line-neutral voltage. Usage of slash rated overcurrent protection devices (OCPDs) is limited to solidly grounded systems according to NEC® Section 240.85. Additionally, corner grounded delta systems require fully rated OCPDs (i.e. full voltage rating rather than slash voltage rating).

For DC, it is possible to encounter slash rating for 3-wire products where the larger number represents the line-line voltage and the smaller number represents the line-midpoint (or equivalently line-neutral) voltage (see Figure 5.3). NEC® Section 240.85 applies here as well. Note that voltage ratings with a “slash” are sometimes marketed for DC chargers, e.g. 250/260 VDC but this represents nominal/float voltage. Similarly, circuit breakers evaluated to UL 489, Annex I for use with battery power supplies (BPS) also have a nominal and a maximum dc voltage rating for the same voltage system, and may be marketed with a “slash” in between, e.g. 1250/1500 VDC.

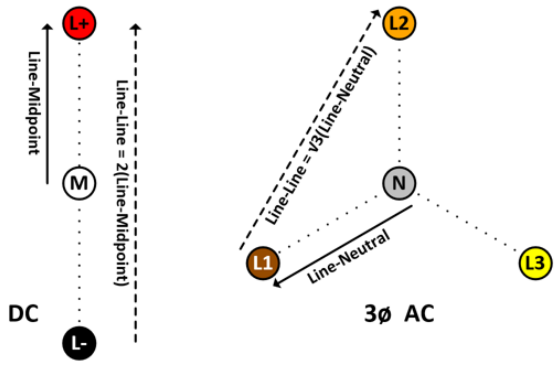


Figure 5.3. DC Midpoint and AC Neutral (Ref.: Current/OS: DC Distribution for Data Centers)

A set of example ratings is supplied in Table 5.3 for AC 1-phase (e.g. 120/240 VAC) and 3-phase (e.g. 480Y/277 VAC) as well as DC 3-wire (e.g. 125/250 VDC and 400/800 VDC) voltages.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

Table 5.3. Example Slash Voltage Ratings

Rating	Line-Neutral (Line-Midpoint)	Line-Line	NEC® Grounding Type Compatibility	
			Solid Grounded (TN-S)	Impedance Grounded or Ungrounded (IT)
480Y/277 VAC	277 VAC	480 VAC	Yes	No
480 VAC	480 VAC	480 VAC	Yes	Yes
120/240 VAC	120 VAC	240 VAC	Yes	No
240 VAC	240 VAC	240 VAC	Yes	Yes
125/250 VDC	125 VDC	250 VDC	Yes	No
250 VDC	250 VDC	250 VDC	Yes	Yes
350/700 VDC	350 VDC	700 VDC	Yes	No
700 VDC	700 VDC	700 VDC	Yes	Yes
400/800 VDC	400 VDC	800 VDC	Yes	No
800 VDC	800 VDC	800 VDC	Yes	Yes

It is important to note that ±400 VDC implies bipolar symmetrical distribution about the midpoint (neutral). While it is not equivalent to 400/800 VDC, devices which take advantage of the symmetry and lower line-neutral voltage in product testing would be rated as 400/800 VDC as opposed to the full 800 VDC. As such, they would be restricted by the NEC® to usage in solid grounding systems.

5.5.2. Voltage Insulation and Isolation

Looking beyond the topic of slash ratings and design for bipolar symmetry, a closer look at the comparative creepage and clearance requirements for unipolar asymmetry in conjunction with either solid grounding (TN-S) or impedance grounding (IT) is provided in Section 6.3. Additional guidance is also provided in this section for the important additional input factors of Over Voltage level and Pollution Degree.

When transitioning between classifications of MVAC, LVDC, and ELV through conversion, it is critical to consider the need for isolation to avoid the need to carry over higher rated insulation to the new voltage class. In addition, for each separately derived system produced by isolated

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

conversion, it is important to state on the applicable equipment nameplate the grounding system compatibility.

Furthermore, the ground fault detection and/or protection requirements of each system in consideration of the chosen grounding scheme require consideration according to Section 7.

5.6. Voltage Bands in Active DC Microgrids

Passive DC microgrids include simple loads which operate independently, based on the presence or absence of source power within a defined tolerance. Active DC microgrids are capable of coordinating multiple sources, loads, and prosumers, and typically employ voltage droop control of current (or power) for maintaining production and consumption power balance. Droop controlled devices are designed to operate within voltage bands and may also be combined with simple loads, provided the impact to system behavior is limited (as further described in IEC TR 63282 and NPR 9090).

Active DC microgrid voltage bands typically provide greater nominal voltage flexibility than is possible in AC, where frequency, voltage, and phase must remain synchronized with the utility for rotating production sources (according to UL 6200) and inverters (according. to UL 1741). Voltage change over time mirrors the power balance in the active DC microgrid, i.e., if more power is needed than available, voltage declines and vice-versa. This property can be used for power management, using so-called “droop curves” (see power management) whereby active components (e.g., converters) adjust their power output according to the voltage measured at their terminals.

In addition to the voltage band descriptions found in IEC TR 63282, droop curve descriptions and examples are also available in the system description documents of ODCA and Current/OS based on LVDC commercial and industrial installations.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

5.6.1. Voltage Band Definitions

A summary of basic definitions for voltage bands and associated voltage levels is provided in Table 5.4 as defined by IEC TR 63282. Continuous operation with full functionality is provided within “nominal” voltage band B3.

Temporary overvoltage operation is provided within band B4 (in Current/OS this is presently limited to 10 minutes and in ODCA devices may choose whether to disconnect). Temporary undervoltage operation is provided in band B2 (in Current/OS this is limited to critical loads and in ODCA critical loads may operate together with derated non-critical loads).
Above the overvoltage band B4 and below the undervoltage band B2, the system operates for a very limited time. Examples are short voltage spikes (e.g. from switching operations) or short voltage dips (e.g. for ride-through capability or prior to selective interruption of a fault condition).

It’s helpful to consider the distinction between the operating voltage range for overvoltage protection devices (e.g. surge protection devices) as voltage band B6 from the operating voltage range for overvoltage and overcurrent protection devices (e.g. solid state circuit breaker SSCB) as B2 through B5 in Current/OS and B2 through B4 in ODCA.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

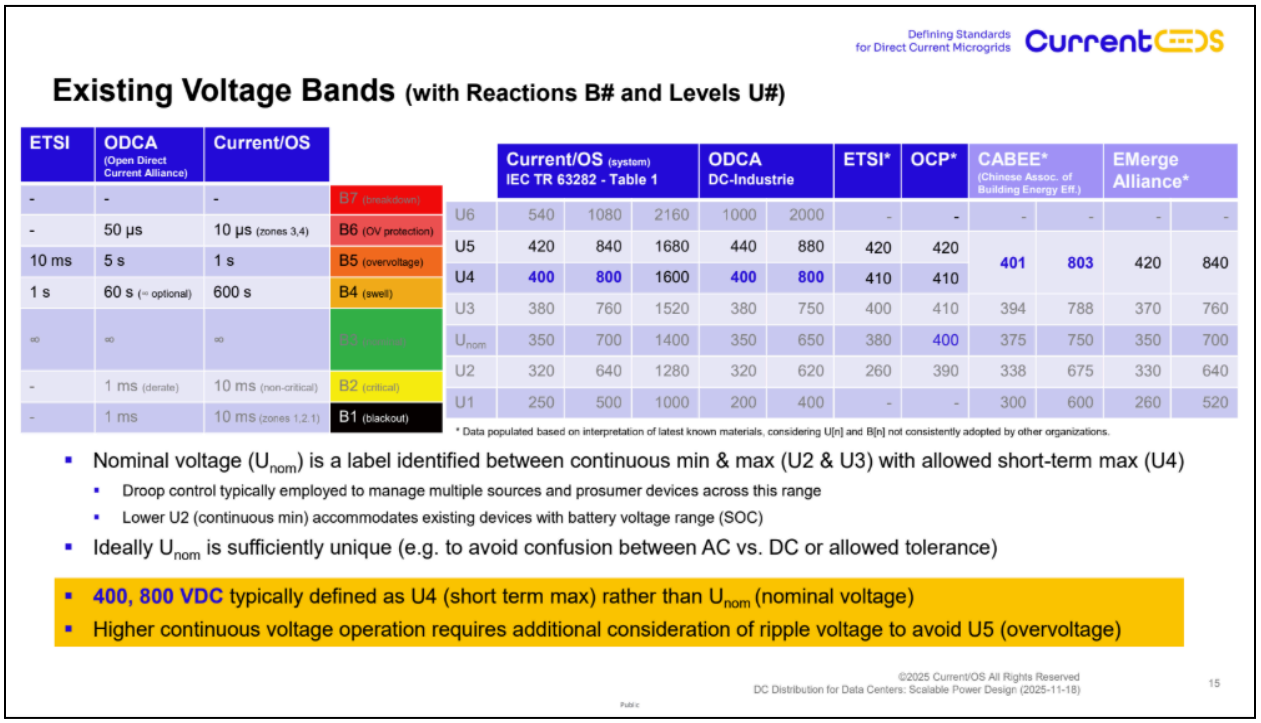
Table 5.4. Voltage Band and Level Definitions as defined by IEC TR 63282

Band	Short Description	Level	Short Description
B7	Prohibited band (U_6 … ∞)		
		U_6	Maximum voltage that equipment overvoltage protection devices clamp
B6	Overvoltage protection devices operation band (U_5 … U_6)		
		U_5	Overvoltage disconnecting voltage
B5	Overvoltage protection devices non-operating band (U_4 … U_5)		
		U_4	Maximum temporary operating voltage
B4	Switching, commutation and protection devices operation band (U_3 … U_4)		
		U_3	Maximum continuous operating voltage
B3	Nominal band (U_2 … U_3)		
		U_2	Minimum voltage for continuous operation with full functionality
B2	Critical band (U1 … U2)		
		U_1	Undervoltage disconnecting voltage
B1	Blackout band (0 ... U1)		

5.6.2. Existing Voltage Bands

An overview of existing voltage band thresholds is provided in Figure 5.4 for DC microgrids (as defined by Current/OS, ODCA, CABEE, and EMerge Alliance) along with data centers (as defined by ETSI and OCP). Note that OCP values are identified based on Mt. Diablo and HPR v4 100kW power shelf specifications as shown in Figure 5.5.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.



Defining Standards for Direct Current Microgrids Current/OS

Existing Voltage Bands (with Reactions B# and Levels U#)

ETSI	ODCA (Open Direct Current Alliance)	Current/OS	
-	-	-	B7 (breakdown)
-	50 µs	10 µs (zones 3,4)	B6 (OV protection)
10 ms	5 s	1 s	B5 (overvoltage)
1 s	60 s (∞ optional)	600 s	B4 (swell)
∞	∞	∞	B3 (normal)
-	1 ms (derate)	10 ms (non-critical)	B2 (critical)
-	1 ms	10 ms (zones 1,2,1)	B1 (blackout)

	Current/OS (system) IEC TR 63282 - Table 1			ODCA DC-Industrie		ETSI*	OCP*	CABEE* (Chinese Assoc. of Building Energy Eff.)		EMerge Alliance*	
U6	540	1080	2160	1000	2000	-	-	-	-	-	-
U5	420	840	1680	440	880	420	420	401	803	420	840
U4	400	800	1600	400	800	410	410				
U3	380	760	1520	380	750	400	410	394	788	370	760
U_{nom}	350	700	1400	350	650	380	400	375	750	350	700
U2	320	640	1280	320	620	260	390	338	675	330	640
U1	250	500	1000	200	400	-	-	300	600	260	520

* Data populated based on interpretation of latest known materials, considering U[n] and B[n] not consistently adopted by other organizations.

- Nominal voltage (U_{nom}) is a label identified between continuous min & max (U2 & U3) with allowed short-term max (U4)
  - Droop control typically employed to manage multiple sources and prosumer devices across this range
  - Lower U2 (continuous min) accommodates existing devices with battery voltage range (SOC)
- Ideally U_{nom} is sufficiently unique (e.g. to avoid confusion between AC vs. DC or allowed tolerance)
- **400, 800 VDC** typically defined as U4 (short term max) rather than U_{nom} (nominal voltage)
- Higher continuous voltage operation requires additional consideration of ripple voltage to avoid U5 (overvoltage)

©2025 Current/OS All Rights Reserved
DC Distribution for Data Centers: Scalable Power Design (2025-11-18) 15

Figure 5.4. Existing Voltage Bands (Ref.: Current/OS: DC Distribution for Data Centers)

Note that Current/OS, EMerge Alliance, and ODCA (recent update from 650V shown in Figure 5.4) have all adopted the 700V and 350V nominal voltages (with 1400V additionally defined by Current/OS). This aligns with the commonly used nominal voltages defined in IEC TR 63282 and soon to be published in an updated version of IEC 60038.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

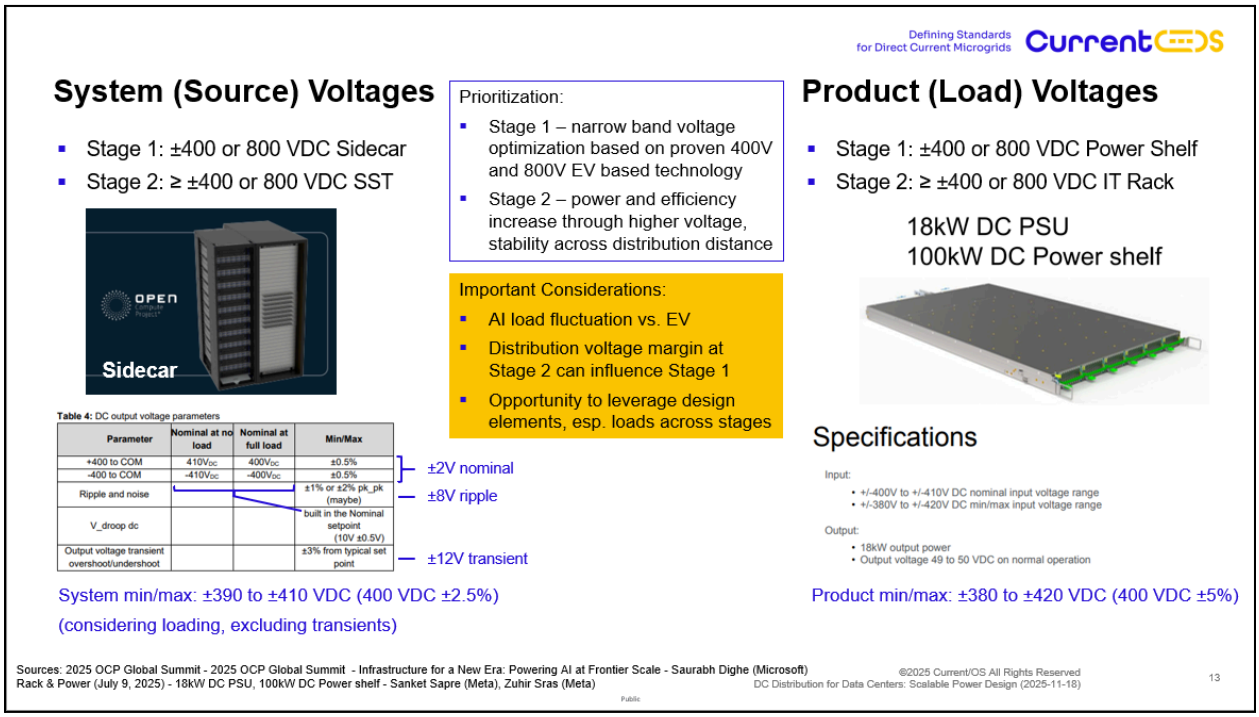


Figure 5.5. Basis for OCP Voltage Bands (Ref.: Current/OS: DC Distribution for Data Centers)

It is important to observe that historically for DC microgrids, operation up to 800 VDC has been limited to temporary overvoltage band B4 (U4 level) with margin above in voltage band B5 (U5 level) to avoid conduction by surge protective devices (e.g. TVSs) as shown by the yellow curve in Figure 5.6.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

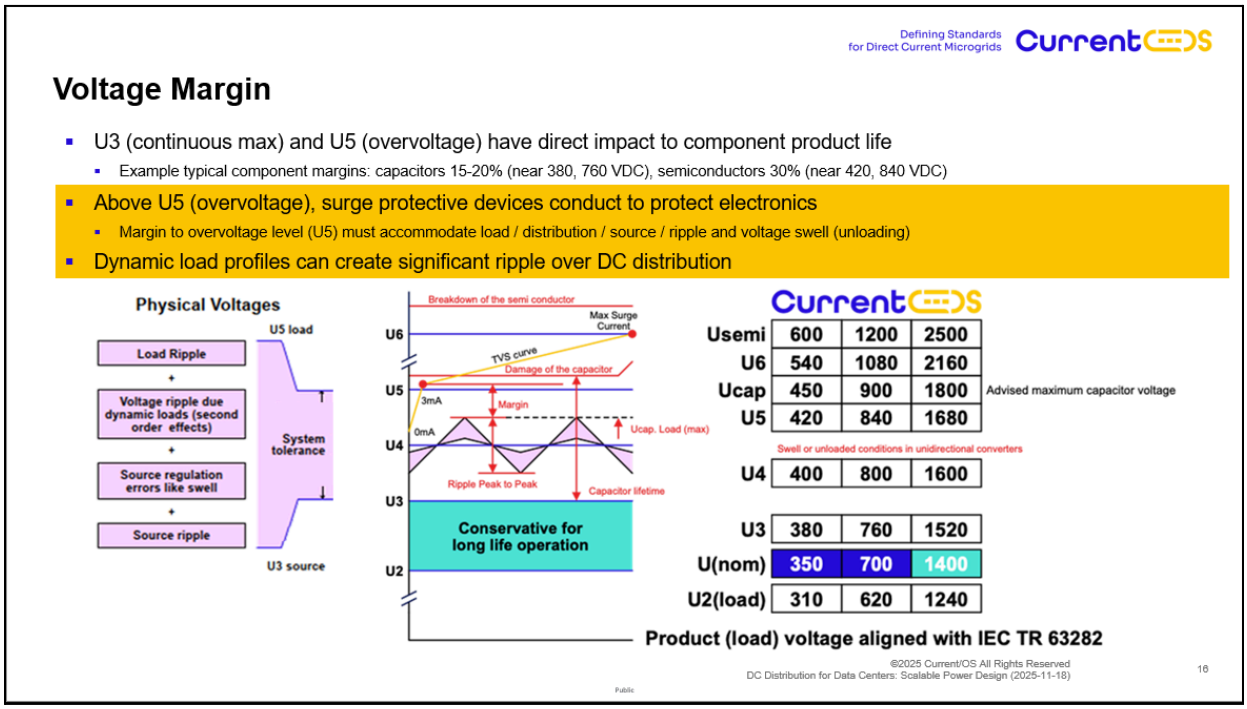


Figure 5.6. Surge Protective Device Margin (Ref.: Current/OS: DC Distribution for Data Centers)

5.6.3. Narrow vs. Wide Voltage Bands

Traditionally, DC microgrid design has embraced wide voltage bands with ample overvoltage margin in band B5 as a means of facilitating “pre-designed” DC systems constructed from interoperable “off-the-shelf” components to provide interoperability among manufacturers, optimize equipment usable life, and minimize custom engineering.

A wide operation range of voltage bands B3 to B4 for normal devices and including B2 for critical loads, has the additional application benefit of minimizing round-trip energy storage losses by aligning with a typical state-of-charge range. Additionally, provisioning voltage band B5 accommodates voltage swell associated with unloading of simple loads and regenerative energy recuperation from motor drives.

In data center applications, direct connection to energy storage is unlikely and chilling loads offer natural motor braking due to pump losses associated with circulating coolant. Fan operation is also possible but deceleration would require appropriate control measures to avoid regeneration during deceleration.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

Furthermore, data center applications intend to take advantage of the economies of scale associated with converter designs and semiconductors utilized in electric vehicles where 400V and more recently 800V batteries have become common. However, it is important to note the significant load stability differences between AI compute and EV applications.

Given the focus on optimizing power delivery for compute loads in the compact footprint of the IT rack, ±400 VDC or 800 VDC has become the design point for next generation DC input power shelves receiving power from Mt. Diablo power racks with ±750 VDC or 1500 VDC as a logical progression for fully utilizing the maximum performance possible in LVDC.

Recognizing this need and objective, significant effort has gone into identifying and comparing the power capabilities of DC power shelves to define a framework for DC system interoperability between data centers and DC microgrids. This begins by taking a closer look at the Power Limiting Converters commonly utilized by DC microgrids which feature a built-in capability for derating at lower voltages as shown in Figure 5.7. Here also by slightly increasing the current limit (e.g. increasing from 131.6A to 140.3A), full power for 750 VDC ±5% is possible.

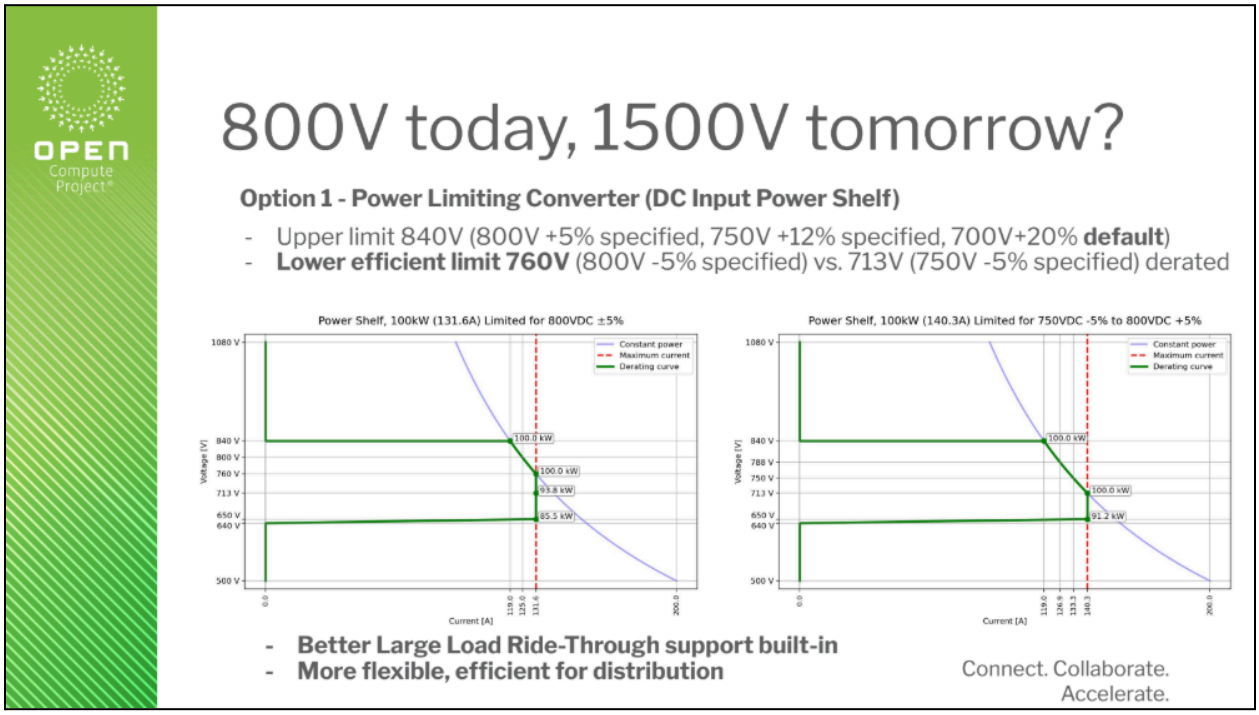


Figure 5.7. Power Limiting Converter (Ref.: OCP: DCF PD Voltage & Power Quality)

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

By contrast, another popular converter design being incorporated into data center power shelves is that of the Fixed Ratio with Point-of-Load Limiting Converter. This type of converter does not have the capability to enforce a fixed current limit, and as voltage decreases, it draws more current and employs protection to limit the duration of undervoltage operation as shown in Figure 5.8. Note that operation at lower voltage range is possible, but requires firmware modification with corresponding impact of derated performance across the entire curve.

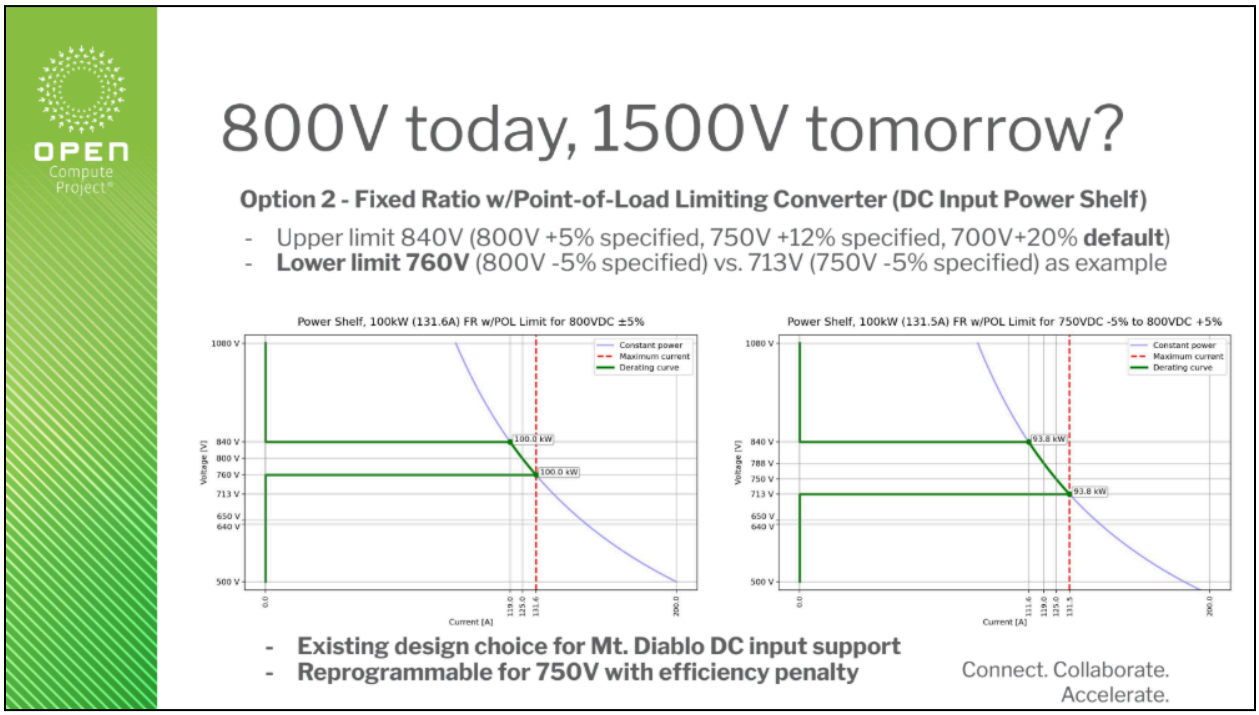


Figure 5.8. Power Limiting Converter (Ref.: OCP: DCF PD Voltage & Power Quality)

Ultimately the advantages of both narrowband DC for compute load space optimization and wideband DC for power distribution margin flexibility are possible to be jointly realized by means of pairing 1500 or ±750 VDC wideband power distribution with 800 VDC narrow band rack utilization (Scalability for Option 1, Figure 5.9). Alternatively, when optimizing for efficiency, usage ±750 VDC within the rack directly is also possible (Scalability for Option 2, Figure 5.9).

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

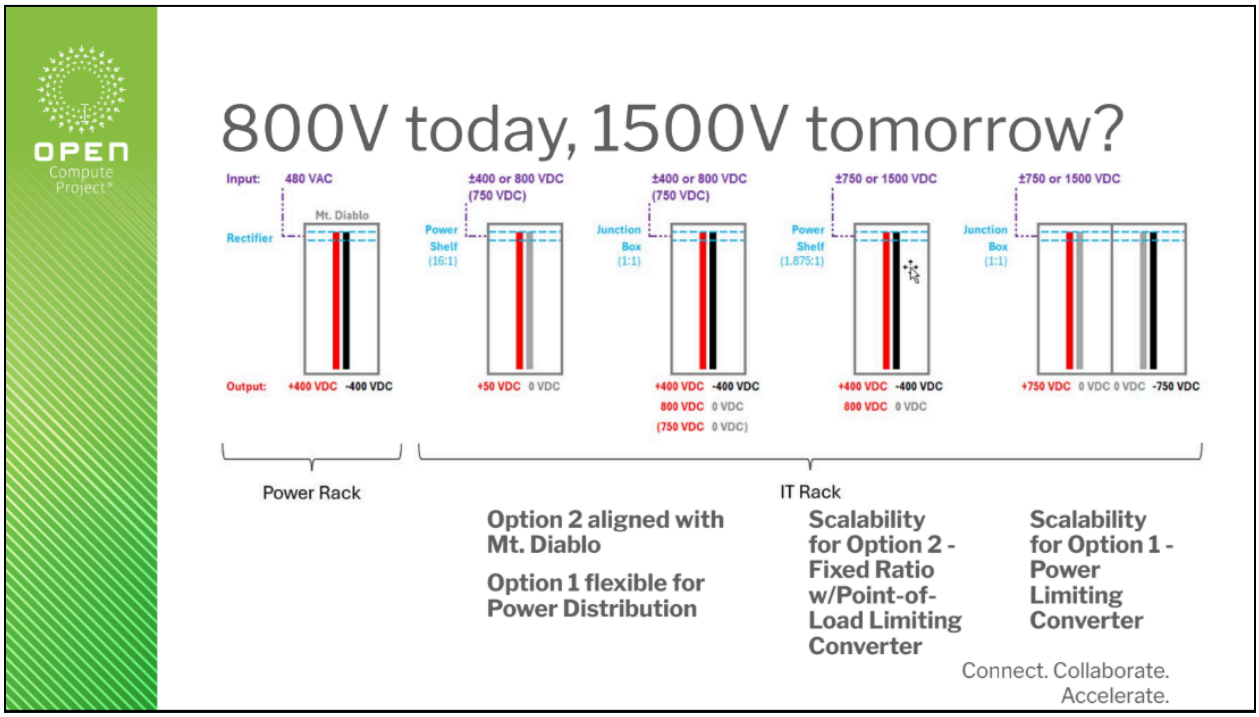


Figure 5.9. Scalability Paths (Ref.: OCP: DCF PD Voltage & Power Quality)

5.7. Interoperable Voltage Band Proposal

The challenge of delivering DC power to IT racks begins with the Mt. Diablo side car concept where interconnect cables may be only 1 meter long to supply the IT rack. DC power distribution delivery from the gray space to the white space is a next level challenge as factors like line inductance for 10's or even 100's of meters become important design considerations factoring in the cyclic load variability of AI compute.

In the interest of building on over a decade of DC power distribution design experience captured in specifications developed by Current/OS and ODCA as well as advancing an ecosystem where economies of scale can be leveraged across data centers and DC microgrids, the following are a set of principles for harmonization of critical design aspects considering present and future data center needs.

The first principle is to consider the distinction between firmware adjustable values and hardware anchor points of ecosystem design. Toward this end, we consider the U2 and U5 levels as the initial boundaries constraining operation in voltage bands B3 to B5. Considering further that droop curves can be used to configure higher voltage limits than U2 without

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

detrimental effect (lower startup voltage than U2 would require design changes to existing equipment to handle higher inrush currents), we are left with the key interoperability anchor point of U5 (840 VDC) as shown in Figure 5.10.

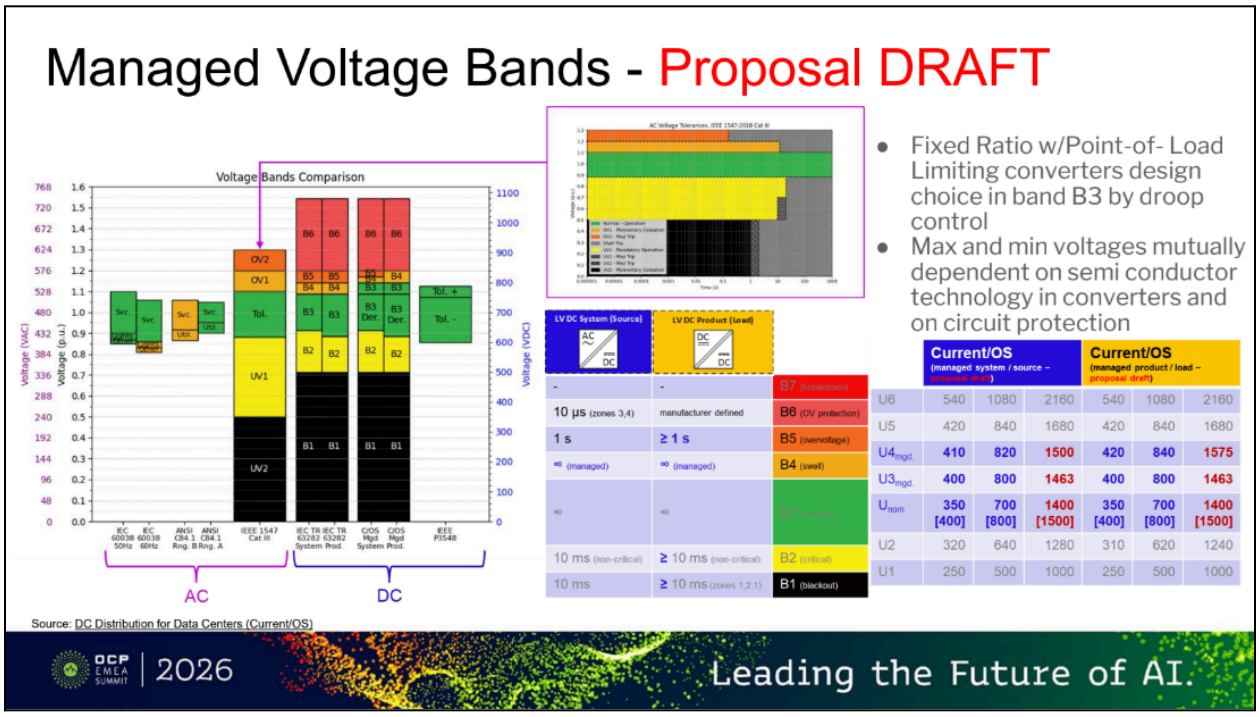


Figure 5.10. Managed Voltage Bands (Ref.: Current/OS: Insights from OCP EMEA 2026)

The second principle additional to this is that we can consider a portion of the B3 band below 800VDC -5% (760 VDC) as optional, thereby accommodating Option 2 Fixed Ratio with Power Limiting converters.

Ultimately, this leads to consideration of the specifics regarding implementation of a narrow band voltage curve. Here it is important to recognize that there is prior art which can be referenced from the 2011 timeframe when 380 VDC was being proposed by EPRI, EMerge Alliance, ETSI, and others as shown in Figure 5.11.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

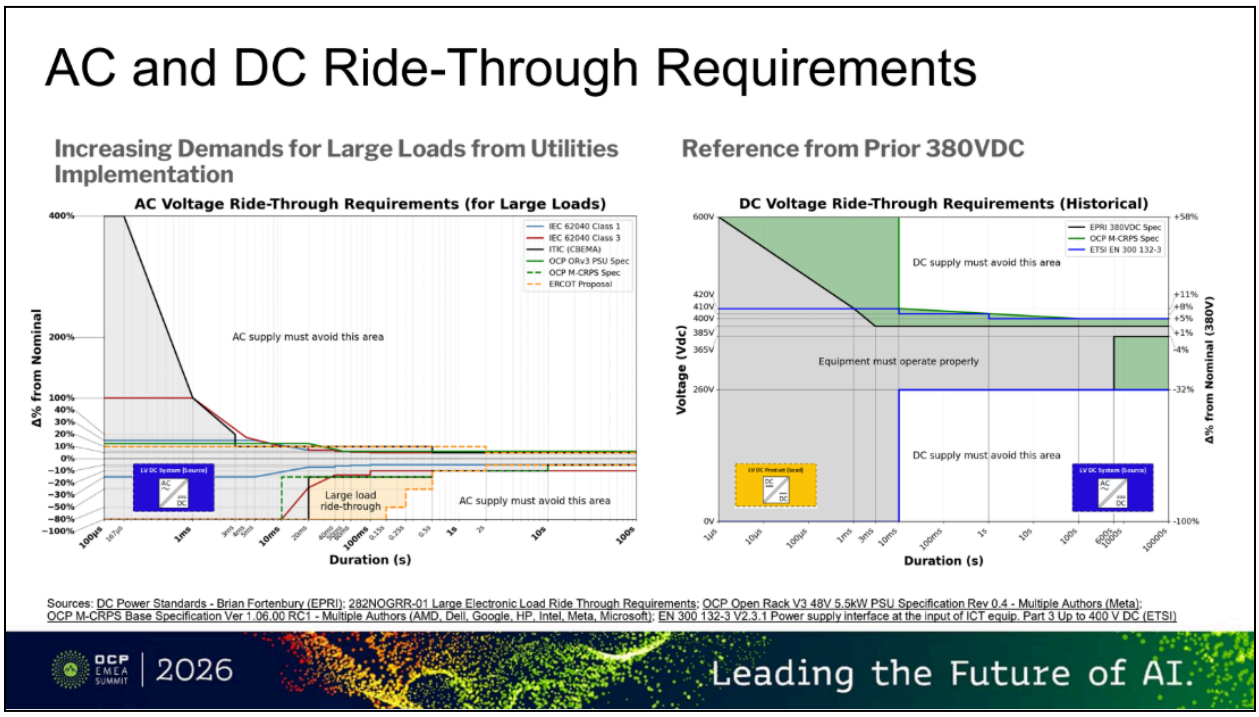


Figure 5.11. AC and Prior DC Ride-Through (Ref.: Current/OS: Insights from OCP EMEA 2026)

Since that time, voltage ride through requirements for large loads have become an additional consideration, driving a closer look at the undervoltage portion of the DC equivalent to the AC ITIC curve. It is important to recognize that while the DC and AC undervoltage ride-through regions share an important connection, fulfillment of DC undervoltage ride-through alone with energy storage does not satisfy the need to continue drawing power from the utility to support the grid. This could mean short-time operation at half power at 50% voltage for example with additional hardware support, as shown in Figure 5.12. The needs of the AC utility (see Section 5.8) are distinct from and do not necessarily translate to the needs of LVDC power distribution. A further exploration of those LVDC power distribution needs is expected to follow the release of a draft OCP SST source specification.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

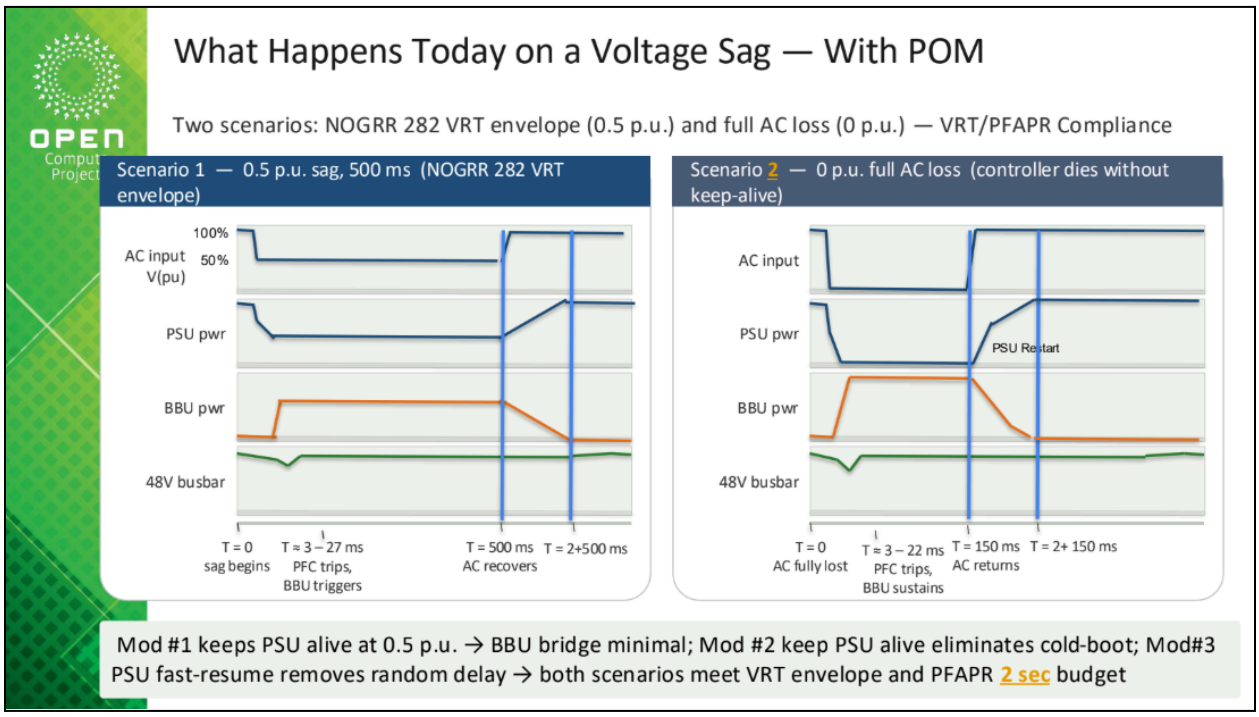


Figure 5.12. AC Ride Through Enablement (Ref.: Power Orchestration Module Proposal)

This undervoltage portion of the DC voltage bands definition is particularly relevant to load converter (power shelf) hold-up capability as additional energy storage outside the IT rack may be required to ensure continued operation. Meanwhile, the overvoltage portion of the DC voltage bands definition is particularly relevant to protection and here, as mentioned previously in Section 5.6.1, U5 (840 VDC) serves as the highest designed protection voltage value according to Current/OS (see Figure 5.14). For this reason, it is imperative that sources are constrained to not produce a higher voltage. Additionally, there is a pending interpretive guidance from IEC TC108 requested to clarify existing IEC 62368-1 600V scope as permitting 850 VDC. This may come sooner than a proposed modification in scope to permit 1000 VAC / 1500 VDC.

Generalized existing voltage-time curves as implemented in Current/OS and ODCA system specifications are shown in Figure 5.13 along with an initial specialized proposal for data centers from Current/OS as shown in Figure 5.14.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

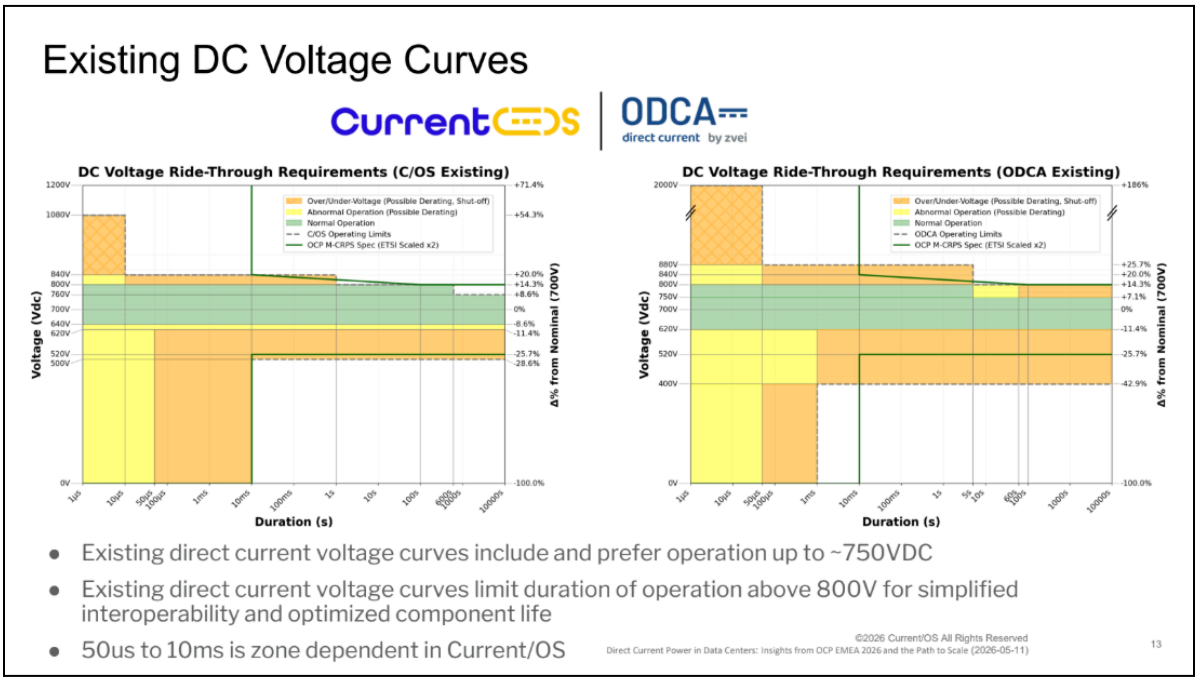


Figure 5.13. Existing Current/OS & ODCA Voltage-Time Specifications

(Ref.: Current/OS: Insights from OCP EMEA 2026)

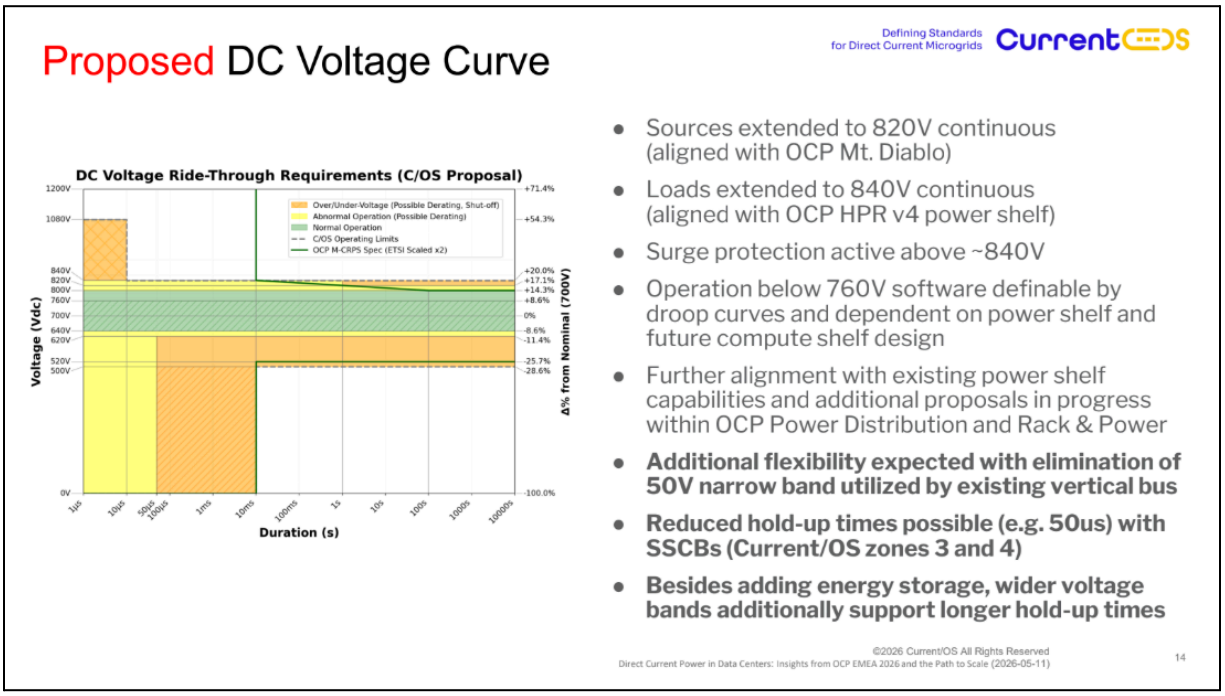


Figure 5.14. Initial Current/OS Proposed Voltage-Time Specification (Draft)

(Ref.: Current/OS: Insights from OCP EMEA 2026)

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

A second proposal based on power rack (side car) and IT rack capabilities of the LVDC input power shelf and future compute shelf has been developed (see Figure 5.15) through coordination between the joint workstream of OCP DCF Power Distribution and Rack & Power (see also Section 9.4).

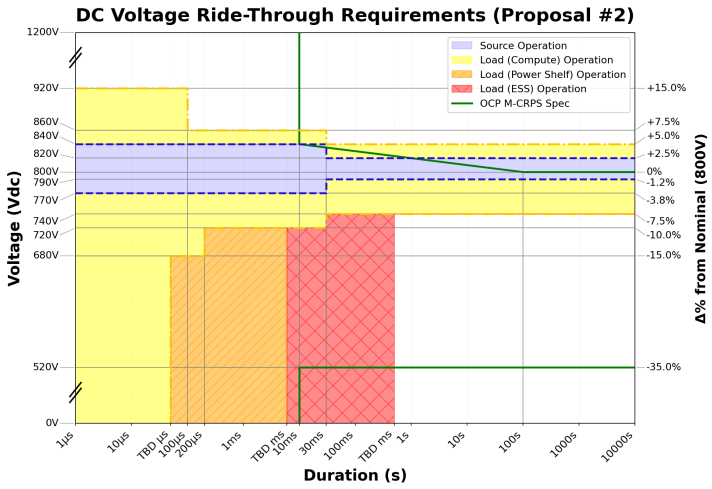


Figure 5.15. Subsequent Joint DCF PD / R&P Proposed Voltage-Time Specification (Draft)

Further alignment and refinement of both proposals is ongoing (particularly in regard to allowed vs. required operation above 840 VDC and required short-term withstand operation to 1080V) and is planned to continue into a future edition of this white paper after publication of the OCP SST draft specification, as potential impact to existing semiconductor technology classes must be considered.

5.7.1. Low Voltage Limit Considerations

As can be seen in the definitions provided in Section 5.2, the limit for low voltage is 1500 VDC. However, an important nuance is that IEC TS 63527 defines this limit in reference to the average voltage whereas the other definitions provided define in the context of nominal voltage which does not have to be the same as the average voltage.

A further consideration is that while ANSI C84.1 does not define an analogous nominal voltage of 1000 VAC, the 5.08% tolerance it commonly applies to other nominal voltages such as 600V results in a value of 1058 VAC which aligns to the maximum voltage defined in ANSI C37.50 for power breakers.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

Considering further on this, it is not uncommon to utilize equipment rated to 1000 VAC in a lower voltage, e.g. 600 VAC system. The issue of name plate labels should then be differentiated for systems / sources as compared with products / loads. An interesting consideration in regard to sources is that whereas Distributed Energy Resources (DERs) are required to accommodate continuous 110% overvoltage, they are typically utility grid following and thus synchronize with the supply of the utility rather than defining their own voltage level. However, notably this same 110% continuous overvoltage level is carried forward to grid forming operation when islanded for DER systems according to UL 3001.

ANSI C84.1 incorporates a factor of 5% for line losses and utilization voltage. Considering that the same line losses exist between a DER source point of connection (inverter terminals) and the point of common coupling with utility (at the meter), it is reasonable to consider that a source tolerance of at least 5% is generally possible – so the question comes down to what to do at the limit.

Ultimately, three likely sufficient but not necessarily required rules are as follows:

1) Avoid setpoints above 1500 VDC. Note that this has an important bearing on series usage of 800V sources. Due to the possibility of 1600V operation, max voltage limit setpoint restriction will be necessary, e.g. at 750V or 1500V. There is analogous max current limit setpoint restriction defined in context of conductor / busbar overload control via Power Control Systems compliant with UL 3141 Section 6.16 and NEC® Article 130 Part II.
2) Avoid average above 1500 VDC (as per IEC TS 63527). This may come with a caveat to be clarified regarding open terminal floating voltage as opposed to loaded condition.
3) Avoid extended excursions above 1500 VDC. Additional work is needed to define a proposal here on the basis of transient and swell definitions.

In addition to the line-line voltage constraint of 1000 VAC / 1500 VDC, an additional corollary may apply in certain locations. For example, in Britain (e.g. BS 7671) there is an additional definitional requirement from BS 7671 for line-earth voltage being constrained to 600 VAC / 900 VDC.

5.7.2. Voltage Band Impact of Static Line Losses

Current flowing through a distribution cable or bus bar causes voltage drop due to conductor resistance. Given the large physical dimensions of data center facilities (and long corresponding

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

possible distribution distances), this voltage drop can be significant and needs to be considered given the finite width of voltage band B3 (see Figure 5.4).

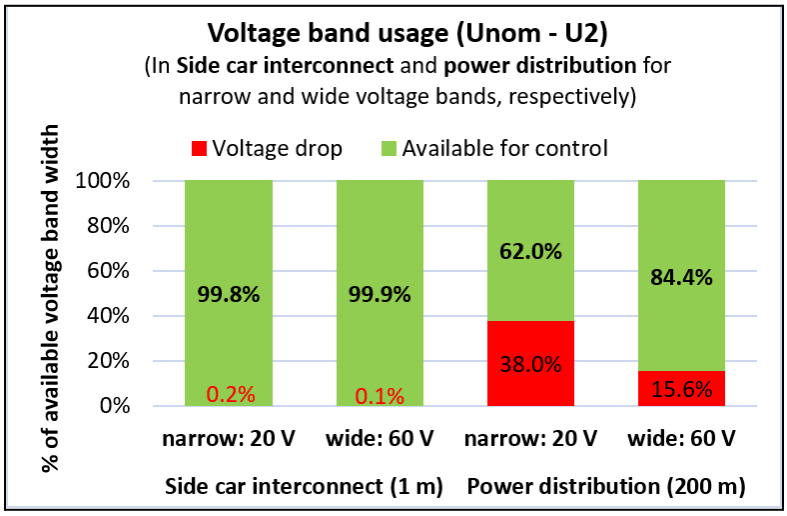


Figure 5.16 Voltage Drop as Percent of Available Voltage Band Width

Figure 5.16 compares the effect of a narrow voltage band (e.g. 40VDC) with a larger voltage band (e.g. 120VDC) for a side car interconnect (e.g. 1m distance) and power distribution (e.g. 200m distance).

For each option, power is 200kW (250A at 800VDC) with minimum required copper cross section of 120 mm² as per IEC 60364-5-52, tables B52.1 and B52.2 (single isolated cables in a conduit). A larger cross section reduces voltage drop (and power loss) but the relative difference for narrow and wider voltage bands, respectively, remains unchanged. The voltage band width in this example is the difference between the nominal voltage (e.g. U_{nom} = 800VDC) and the lower end voltage of the nominal voltage band B3 (e.g. U2 = 780VDC). For the short interconnect distance from side car to IT rack, voltage drop is not relevant (green “available for control” portion dominant with >99%) while for the power distribution along the data hall, the “available for control” portion would only be 12.4VDC out of 20VDC in the narrow band case.

In this example, dynamic reactive line loss effects were neglected and only resistive line loss effects were considered. Even so, this demonstrates the benefits of a wider voltage band for longer distance distribution to enable a wider range of voltage for more stable control and less precise measurement while managing power balance by droop curves (refer to Section 5.8.1).

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

5.8. Power Quality and Grid Interaction

Data centers using LVDC power distribution must maintain acceptable internal DC bus voltage quality (voltage-time behavior within voltage bands) while also behaving as a stable and predictable load at the utility point of common coupling (PCC). In this section, power quality is considered in a broad system architecture sense. It includes DC bus power balancing, the propagation of upstream utility AC voltage and frequency disturbances to the DC system, and the active power response of the facility during and after those disturbances.

5.8.1. AC Power Quality at the Utility Interface

Whereas utility grid support is commonly required for bidirectional distributed energy resources (DERs, e.g. IEEE 1547) and inverter based resources (IBRs, e.g. IEEE 2800) in the form of both ride-through and active / reactive power support, increasingly utilities are additionally defining ride-through requirements for unidirectional large loads as shown in Figure 5.17 (latest ERCOT info compared with Figure 5.11).

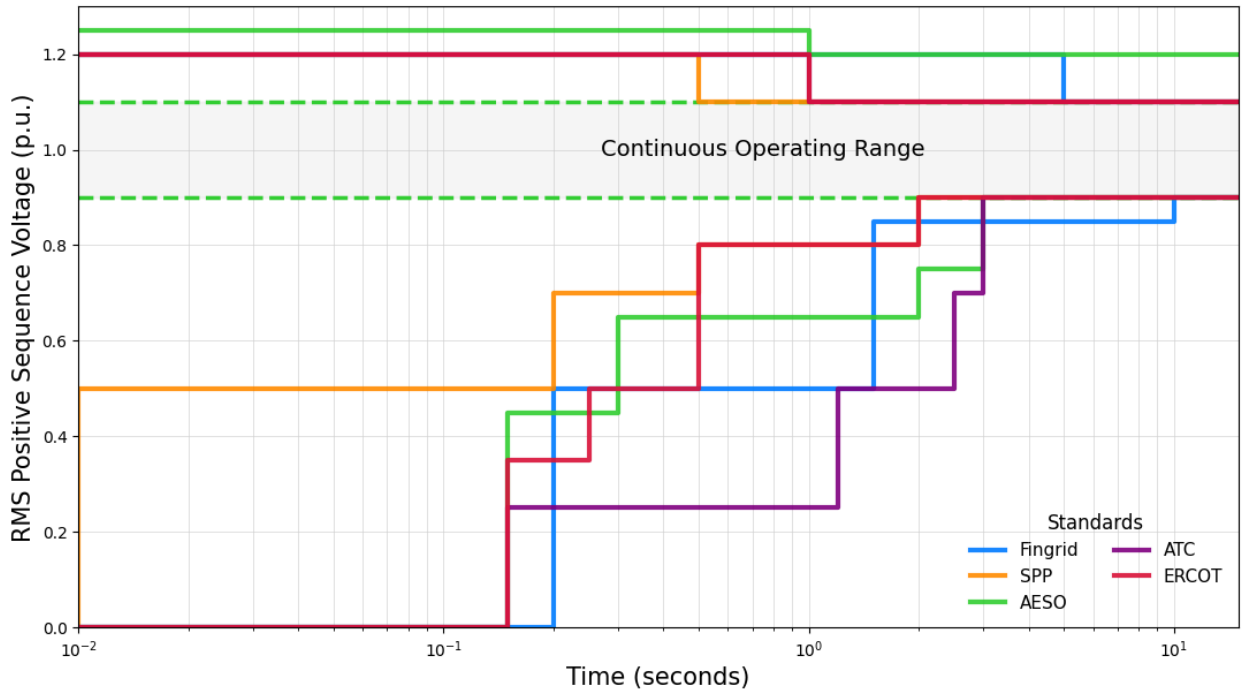


Figure 5.17. Proposed Large Load Voltage Ride-Through Requirements

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

Increasingly, stability interaction between the grid and the data center has bidirectional impact regardless of whether power flow is bidirectional. An upstream grid disturbance can impact rectifiers, power supplies, energy storage systems, and protection controls. Conversely, the simultaneous transfer or trip of large power electronic loads can amplify the disturbance seen by the grid.

On July 10, 2024, a 230 kV transmission line fault and repeated automatic reclosing in the Eastern Interconnection coincided with customer initiated loss of approximately 1500 MW of data center type load. The North American Electric Reliability Corporation (NERC) reported that frequency increased to 60.047 Hz and voltage to approximately 1.07 per unit (p.u.) following the load loss [10] as shown in Figure 5.18.

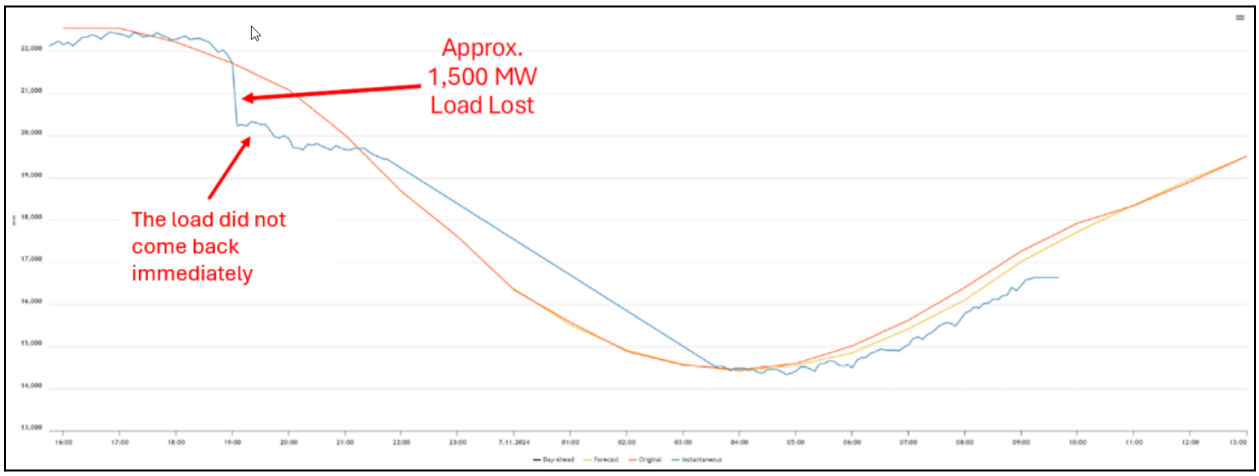


Figure 5.18. Loss of ~1500 MW of Data Center Type Load (Ref.: NERC Incident Review [10])

Additional discussion of changing AC power quality requirements for sustained power draw from the utility in case of disturbances and for recovery after fault clearing is covered in Section 8.1.

5.8.2. DC Power Management with Droop Curves

In DC systems, an imbalance in power supply and demand leads to a change in voltage that is measurable at any point in the installation. Higher demand than supply leads to lower voltage and vice-versa. Since DC systems do not have to stick to a fixed voltage value (a significant advantage over rigid voltage-frequency demands of AC systems), this behavior is used to

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

balance power among sources, loads, and prosumers (e.g. storage) without the need for a high-speed central controller.

Active components such as DC-DC converters measure the voltage at their terminals and adjust their current output according to pre-defined current-voltage curves, known as droop curves. Figure 5.19 shows an example of droop curves for three sources: infeed from the grid (red), photovoltaic (green), and energy storage (blue). In addition to regulating the power supply, loads can also be adjusted – preferably those with a larger time constant, such as heating or cooling units for the building. Portions of curves below the zero current line indicate current drawn from a source while portions of curves above the zero current line indicate current feeding into storage or the supply grid.

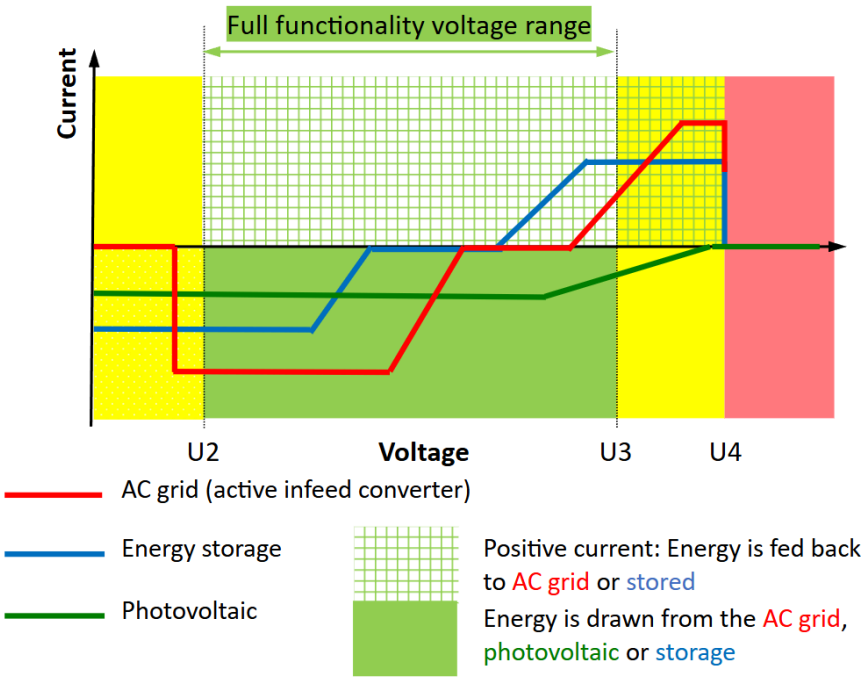


Figure 5.19. Example Droop Curve Power Management (Ref.: System Description [ODCA])

Further discussion of DC power quality requirements and their interaction with AC power quality requirements is planned in a future edition of this white paper after publication of the OCP SST draft specification.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

6. Polarities & System Earthing

This chapter defines and evaluates polarity configurations and earthing schemes used in direct current (DC) data center power systems. Its purpose is to clarify possible design options, identify technical tradeoffs, and support consistent implementation and commonalities across different system architectures.

The chapter reviews unipolar and bipolar systems, including variants with fully distributed, partially distributed, or non-distributed midpoints (M). These configurations are further reviewed in conjunction with earthing approaches TN-S and IT, including the use of single and multiple grounding points. As detailed below in the Section Combinations of Polarities, System Earthing and Voltages, combinations of polarity, earthing scheme, and isolation of the AC/DC interlink are assessed to highlight practical implications, including stray currents and corrosion risks resulting from earthing choices and system topology. Additionally, the interaction between polarity, earthing, and isolation of the AC/DC interlink is analyzed with respect to fault behavior, protection requirements, electrical spacing, personal safety, maintenance and operational considerations.

6.1. Polarities

DC grids can be configured as either bipolar or unipolar. In general, bipolar has two poles, of opposite polarity, that significantly differ from ground while unipolar has one pole that significantly differs with respect to ground. These polarities have different characteristics and can affect the requirements for the technology or connected equipment used in the system.

Table 6.1. Definitions for designations of conductors in a DC grid

Symbol	Definition
L+	Positive Line
M (or N)	Midpoint Conductor (close to ground). Labeled N in NEC®
L-	Negative Line

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

There are different variations of bipolar and unipolar systems that differ based on the presence of the three possible conductors (L+, M and L-) in the whole electrical system or in part of it as shown in the table below:

Table 6.2: Polarities and use of conductors

Variations of Polarities	L+	M (or N)	L-
Bipolar with distributed M	x	x	x
Bipolar with partially distributed M	x	(x)	x
Bipolar without distributed M	x		x
Unipolar (positive)	x	x	
Unipolar (negative)		x	x

The following sections will describe these variants and their main characteristics as they can affect the requirements for the technology or connected equipment used in the system.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

6.1.1. Bipolar

A bipolar system, by definition, has two different polarities towards ground. It consists of at least two distributed conductors, L+ and L-, with opposite polarities relative to a ground. Optionally, a midpoint conductor M, can be included. The M may be connected directly to earth, or tied to earth through an impedance. In the IEC dictionary bipolar systems are defined in entry 601-04-03.

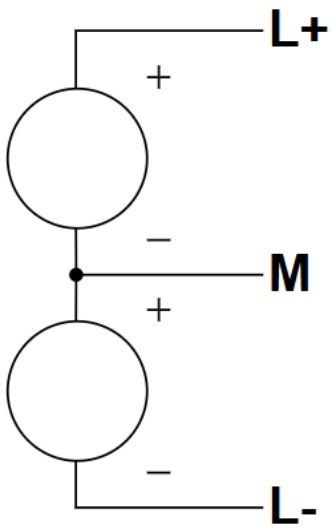


Figure 6.1. Bipolar DC grid circuit model where M connection is optional

Main Characteristics:

- Lower potential difference to earth than unipolar systems, leading to smaller isolation distance and device size.
- Higher power transfer capability compared to unipolar equivalent for a given voltage-to-ground potential
- At least 2 poles have to be protected (3 poles if IT grounding)
- Lower insulation voltage line-to-earth compared to line-to-line
- Can be compatible with non-isolated AC/DC interlink implementation if TN-S grounded on the AC side

There are multiple M conductor distribution options in a bipolar system that are described in the following sections.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

6.1.1.1. Bipolar with Distributed M

This configuration uses a full 3 conductor (plus PE) distribution. It requires all 3 conductors (L+, M, and L-) to be distributed throughout the system and it provides full flexibility to use sources and loads connected in four different configurations:

- L+ and L-
- L+, M and L-
- L+ and M
- M and L-

However, a bipolar system that has sources or loads connected to only one pole (L+ to M, or M to L-), has the potential to become imbalanced if the loading is not well matched. The M conductor could be rated for lower current, and using thinner cross-section, but only if it is ensured that the maximum unbalance is limited or M is appropriately protected against over current. This limits the ability of the system to continue operation on one pole (in the event of a fault in the other pole) to the maximum current rating of the M conductor. The requirement that both poles are required to trip in a M grounded conductor system (as is currently required in NEC® standards) can be configured in some protection devices for Bipolar DC grids.

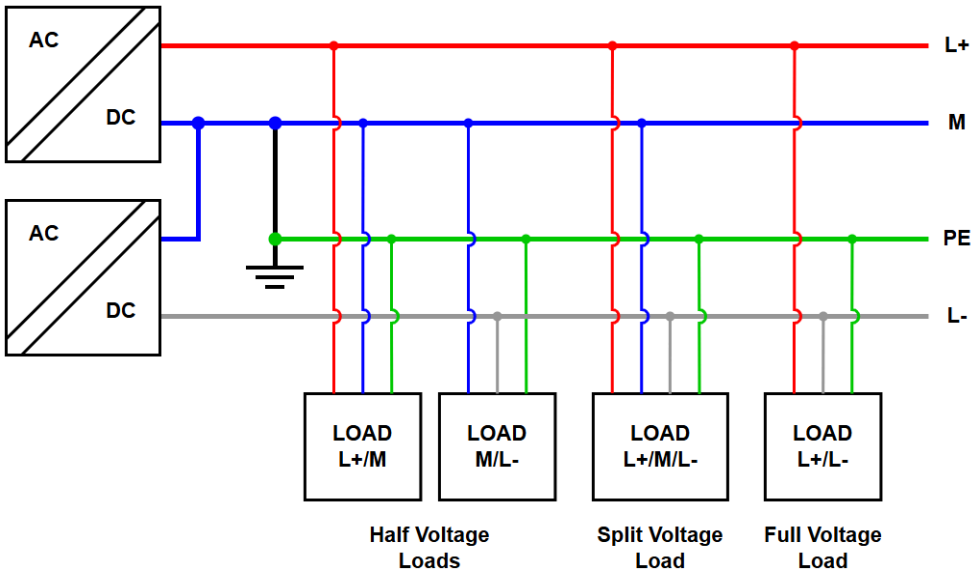


Figure 6.2. Bipolar TN-S grounding with one solid grounding point

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

Main Characteristics:

- Two voltage levels available at any connection point
    - Enables options for 2-level power topologies using 600V or 650V silicon or GaN devices operating on a 350V while the full 700V range may be more suitable for higher power units utilizing SiC devices or multi-level topologies with GaN.
- Unbalanced situations might require additional balancing equipment depending on design and control (if sources feed only L+ and L-)
- Could be designed to continue operation at reduced capacity, if one pole is lost, to enhance availability (provided that sources can handle fully unbalanced load).
- Presence of M conductor throughout the system allows multiple distributed grounding points throughout the system which permits localized ground fault return paths in the system, thereby increasing the system's protective selectivity (more can be seen in Multiple References to Earth).

6.1.1.2. Bipolar with Partially Distributed M

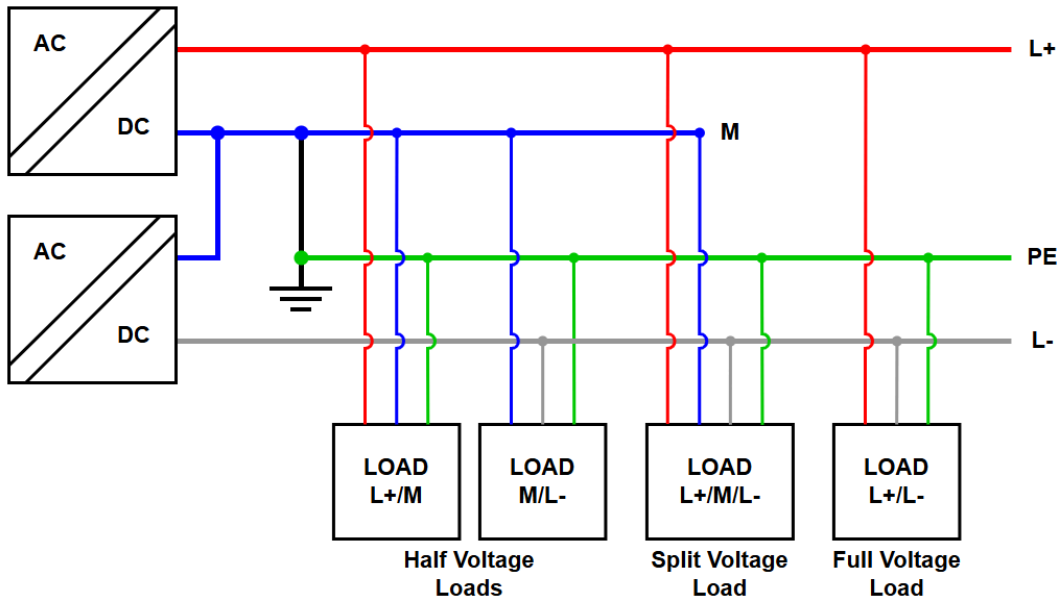


Figure 6.3. Bipolar TN-S grounding with partially distributed M

This option is a system where the M conductor is partially distributed. This allows loads, in some areas of the system, to utilize the M conductor while saving cost by not distributing M to

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

other areas where only full-voltage loads are utilized. (Comparable to 3 phase AC with N for single phase loads and without N for large loads).

Main Characteristics:

- Maintains most characteristics of distributed M
- M conductors can be eliminated and copper saved in circuits that only distribute balanced power.
    - This can apply for circuits that only power large loads connected at L-L voltages
    - This can potentially also be done for longer cables as long as adequate balancing power is available on both ends.

6.1.1.3. Bipolar without Distributed M

This symmetrical system uses two current-carrying active lines only, no midpoint is used. This system provides only one voltage. Note that although the M node is tied to PE at the source, the PE line shall not be used as a return path from loads to M. The PE line is only for protection.

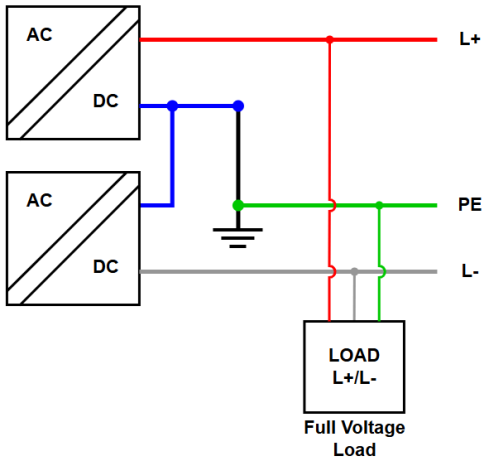


Figure 6.4. Bipolar DC system without distributed M

Main Characteristics:

- Only two active wires needed for the same power
- One available voltage level
- No power balancing needed
- No active wire close to ground can result in lower risk of PE corrosion without residual current monitoring (RCM), see chapter Corrosion.
- Has lower power loss than unbalanced three-wire systems with M

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

6.1.2. Unipolar

A unipolar system consists of one pole and a return conductor, either L+ and M, or M and L-. Same as the bipolar without M, these systems only provide a single voltage level. The M terminal is the ground referenced conductor. A positive system (L+ and M) is generally preferred in a unipolar configuration.

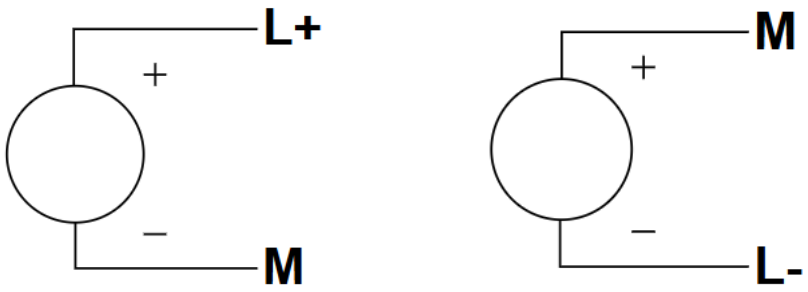


Figure 6.5. A positive unipolar system (Left) and a negative unipolar system (Right)

Since a unipolar system only has 2 active conductors, the distribution network requires the M conductor to be fully routed.

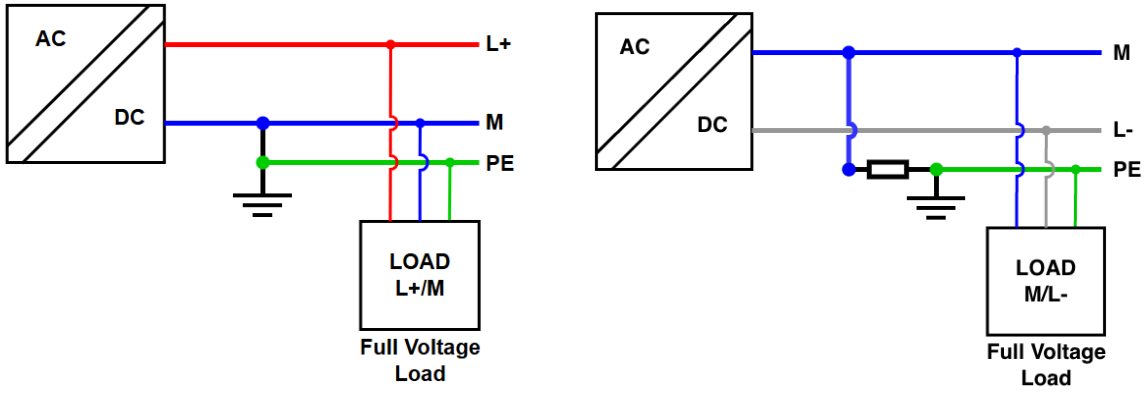


Figure 6.6. TN-S grounded unipolar positive system with L+ and M (Left) and a IT grounded unipolar negative system with M and L- (Right)

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

Main Characteristics:

- Single active conductor relative to a return path
- Only one voltage level available
- Simple protection devices (1 pole protected)
- Reduced power transfer capability relative to a bipolar configuration for a given voltage-to-earth potential.
- Isolation is required if fed from an AC supply. See Figure 6.9 in the isolation section.

6.2. System Earthing

This section begins with two tables intended to clarify some grounding/earthing terminology (and semantic equivalents) taken from NEC® and IEC standards. A third table contains a breakdown of IEC nomenclature used in grounding schemes. The following references used in the tables are listed below:

- NFPA 70 – National Electrical Code (NEC®)
    - NEC® Article 100 + Article 250
- IEC 60050-195 – International Electrotechnical Vocabulary
- IEC 60364 series – Low voltage electrical installations

Table 6.3. Terminology equivalents between UL/NEC® and IEC standards

Concept	UL/NEC (North America)	IEC (International)
General Term	Grounding	Earthing
Resistance or impedance based grounding	High resistance grounding (HRG)/ Low resistance Grounding (LRG)	Impedance earthed
Faults	Ground Fault	Earth Fault
Electrode	Grounding Electrode System	Earth Electrode
Equipment grounding	Equipment grounding conductor (EGC)	System Referencing Conductor (IEC 60364-8-82)

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

Table 6.4. Grounding specific terminology equivalents between UL/NEC and IEC standards

NEC/UL Description	IEC System Type	Explanation / Note
Solidly grounded wye system (AC) or Grounded DC System (DC)	TN-S	Neutral solidly earthed, PE separate
Solidly grounded with combined N/PE	TN-C	Not allowed by NEC in most cases
Ungrounded System	IT System	System isolated from earth or via impedance
Corner-grounded delta	No direct IEC equivalent	North American specific practice
High-resistance grounded	IT or impedance-earthed TN	Depends on implementation

Additionally, Table 6.5 explains these letters and their meanings (IEC 60364-1, ch. 1.9.1).

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

Table 6.5. Earthing Nomenclature: Nomenclature of earthing systems

First letter	Indicates how live parts of the electric system are related to earth:
T	One live conductor is directly connected to earth (terra)
I	All live conductors are isolated from earth
Second letter	Indicates how exposed conductive parts are earthed
N	Exposed conductive parts are connected to the same point where the live parts are connected to earth
T	Exposed conductive parts are connected to an earthing system which is independent of the electrical system
More letters (if any)	Separated by a dash from the first two. Indicates how live parts and earthing conductors are arranged
S	Separate conductors for neutral (midpoint) and protective earth (PE) conductors
C	Combined conductors for midpoint and PE, e.g., PEM or PEL

Additionally, in DC systems, only TN-S or IT systems shall be used. Other earthing schemes such as, TT systems and TN-C systems, are not used because of the risk of corrosion of grounding electrodes or ground-burried metal structures which are significantly higher because earth-fault currents, by design, path through metal (grounding electrode) to electrolyte (earth) interfaces. On the anodic side, the metal corrodes if the potential drop surpasses a (small) voltage threshold.

6.2.1. Isolation of Electrical Systems

Isolation is typically used when transitioning between systems with different voltage levels, safety classes, or grounding requirements. When it comes to grounding, isolation creates two separate systems which can have separate grounding configurations (in the NEC, the new

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

system created by isolation is defined as a separately derived system). On the contrary, no isolation means that you can have a single system with a single grounding configuration. As an example, if a non-isolated interlink (AC/DC) is used, it means that the AC side grounding point could be used as the DC side grounding. This configuration can be seen in Figure 6.8. TN-S-AC.

6.2.1.1. Medium Voltage (MV) to Low Voltage (LV)

Isolation is needed when connecting MVAC or MVDC to LVDC distribution via a SST or dedicated AC transformer. With a dedicated transformer from medium voltage to low voltage, no isolation for the converter is needed (since the transformer does the isolation).

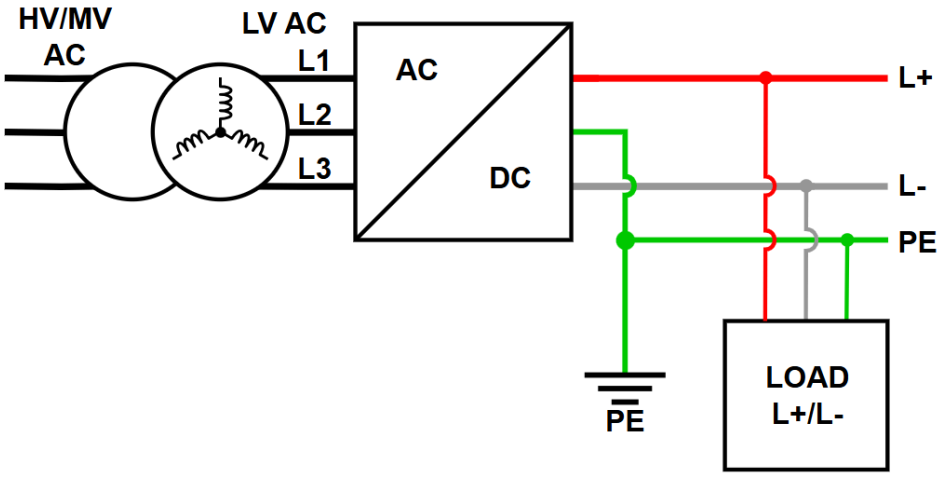


Figure 6.7. Symmetrical power distribution with non-isolated AC/DC converter **with DC side earthing**

6.2.1.2. Low Voltage (LV) to Extra Low Voltage (ELV)

Isolation is needed in ELV applications, for example, in a server or for user-accessible equipment in data centers when a large voltage step is present from LV distribution. This is typically a large voltage step-down where isolated converters are more economically viable, anyways.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

6.2.1.3. Grounded AC side Low Voltage AC (LVAC) to Low Voltage DC (LVDC)

In this scenario, an AC/DC interlink is required to connect to LVDC and it may be either isolated or non-isolated. The following list describes various considerations for isolation and grounding in this case:

- For isolated applications, any of the following polarity and grounding options can be utilized for the LVDC side: unipolar, bipolar, TN-S, IT, etc..
- The AC side hard grounding point can be utilized also as the grounding for a DC grid if a non-isolated interlink is used to interface the two systems. However, if no other AC loads are connected and the LVAC network serves only as an intermediate link, grounding on the AC side provides limited benefit. In such cases, it is preferable to ground the DC side instead.
- For a non-isolated application operating at low-voltage AC, if the AC side is grounded, the DC side must be implemented as a bipolar DC system. A unipolar DC configuration is not suitable because the DC voltage must exceed the peak value of the AC voltage (both positively and negatively) to avoid uncontrolled rectification through the converter.
- Advantages to non-isolated interlink: simpler, lower cost, less bulky, higher efficiency.
- Limitations for non-isolated interlink: Common mode currents (3rd harmonic) → Ac side grounding requires consideration of harmonic effects and DC common mode current restrictions apply as well.

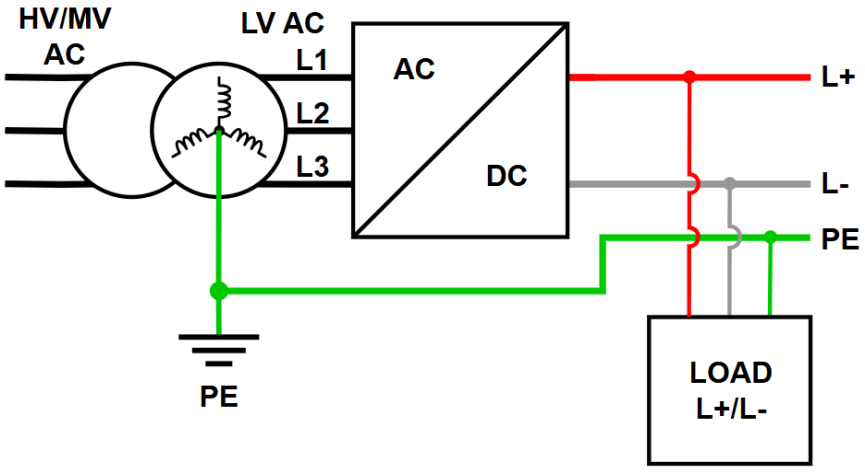


Figure 6.8. Non-isolated-AC-grounded: symmetrical power distribution with non-isolated AC/DC converter **earthed on the AC side**

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

6.2.2. TN-S (Solidly Grounded System)

These systems are characterized by a grounding connection at the source and separate paths for a neutral, current-carrying conductor (if applicable - midpoint M in a DC system) and the protective earthing conductor (PE) that connects all exposed conductive parts bodies in the application with ground. A single solid grounding point in a TN-S system can reduce the potential for DC ground currents and corrosion.

Grounding and connection to the supply grid for a TN-S system with DC-side grounding is shown in Figure 6.9. The AC/DC converter provides a midpoint on the DC side which is grounded

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

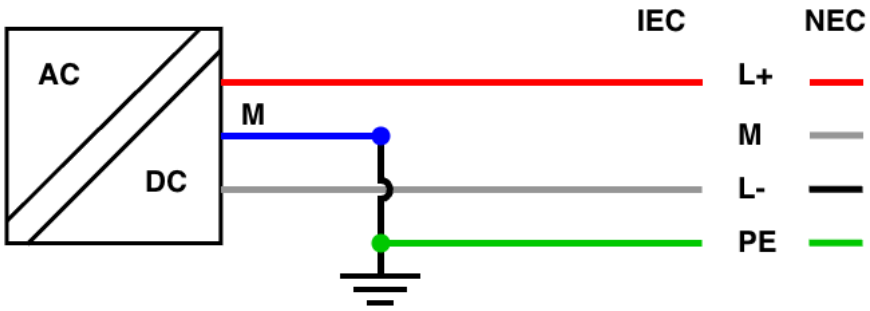


Figure 6.9. TN-S-DC: TN-S system with **earthed on the DC side**

Similarly, the AC-side grounded TN-S system is depicted in Figure 6.8. The connection to ground is most favorably made at the star-point of the AC transformer. A more detailed representation of this is shown in Figure 6.8 above in Section 6.2.1.3.

6.2.2.1. Direct Touch Protection (with RCD)

Additionally, residual current monitoring (RCM) and residual current detection (RCD) can be implemented in a TN-S grounded system. In the case of a bipolar DC grid, the RCD can provide selective protection by clearing an earth fault path on the faulty pole while the other pole remains functional. This is possible if the M conductor is present and fully rated.

Moreover, implementing RCDs at all branch circuits or racks, allows for larger scale systems to be implemented without system level limits on the Y-capacitance to earth as any harmful earth fault currents can be identified and cleared locally by the RCD. This effectively prevents the cumulative Y-capacitance of the entire facility level LVDC grid from discharging at an earth fault in a single circuit. Only the Y-capacitance in the affected circuit will discharge into the earth fault. Tripping curves for DC can be derived from IEC 60479-2 “Effects of current on human beings and livestock”. More information on RCD can be found in Section 7.3.

6.2.2.2. Multiple References to Earth

For DC, multiple solid earthing points or system referencing conductors (SRCs) should be avoided because of the risk of DC currents through earth or other conductive structures that can cause corrosion. Notably, solid grounding at multiple points was a practice in some 50V DC OCP racks and should be avoided in larger DC grids as corrosion would not be limited to the rack but could also affect other parts of the facility.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

On the other hand, multiple *references* to earth can be beneficial for selective protection of the DC grid and can increase resilience and continued partial operation after faults if multiple supplies/sources exist. This can allow for new architectures to emerge which could optimize redundancy in data centers. Additionally, for DC, these additional references to earth can be realized with System Referencing Devices (SRD) that block DC but conduct fault transients for fault protection. (Also commonly named: Capacitive Grounding, PCR (Polarized Cell Replacement), and Capacitor-Diode grounding). SRDs have to be rated for the expected maximum fault current. When using current limiting protection devices, such as semiconductor breakers, this maximum fault current is lower and can make this option more attractive. Typically, they include capacitive behavior in addition to a clamping mechanism to limit the voltage to PE. One last consideration for multiple earth references is for islanding and de-islanding a grid. If multiple SRDs are used throughout the system, the grid can be islanded in multiple configurations as long as an SRD is present. If there are not multiple references to earth, then a Switching Device for Islanding (SDFI) or similar technology could be used to establish an earthing reference for an islanded grid. Otherwise, the electrical system changes to IT and insulation monitoring is required. More information can be found in IEC 60364-8-82.

Multiple references to Earth with SRDs are being discussed for inclusion in the currently ongoing revision of IEC 60364-8-82 on prosumer electrical installations as well as public input for NEC 2029.

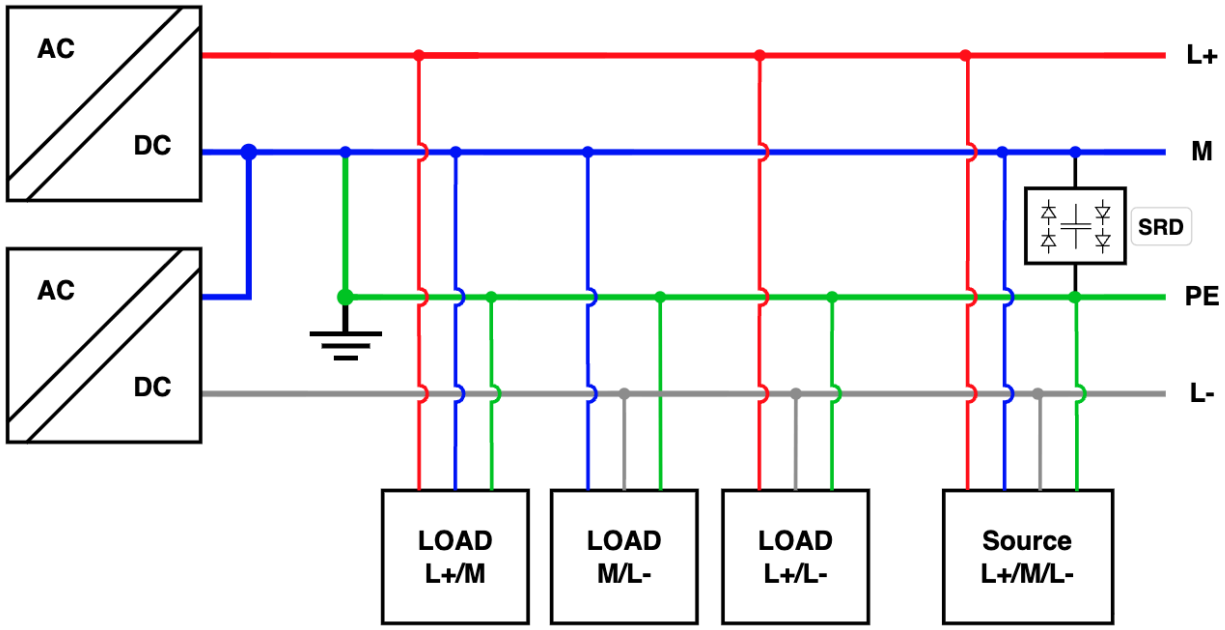


Figure 6.10. Bipolar TN-S with multiple distributed System Referencing Devices (SRDs) added

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

6.2.3. IT (Ungrounded System or High Impedance Grounded)

IT systems are often described as floating systems (ungrounded in the NEC, with no intentional connection to ground) or high impedance grounded systems. In either case, impedance to ground is typically kept higher than 100 Ohm/V for DC or 500 Ohm/V where non-isolated from AC (per UNECE R100, ISO 5474 series, and ISO 6469 series). IT systems are able to continue operation for a certain time after a first ground fault. This enables the operator to either continue services while looking for the cause of the fault, repair it, and then resume normal operation or enable the circuit to perform a controlled safe shutdown instead of an emergency shutdown. Such a system needs components that are designed for the full line-to-line voltage because that voltage appears when one of the active, potential-carrying conductors has that fault-causing connection to ground. Specifically, clearance distances (air gaps) need to be larger than in a grounded TN system, and protection devices could be required in the future to clear a fault, with one pole only, against the full voltage (currently installation standards require overcurrent protection on all poles, but this is not needed, theoretically). Figure 6.11 shows an IT system that has no direct connection to ground on the source (in this case, the AC/DC converter) and it shows isolation monitoring and symmetrization.

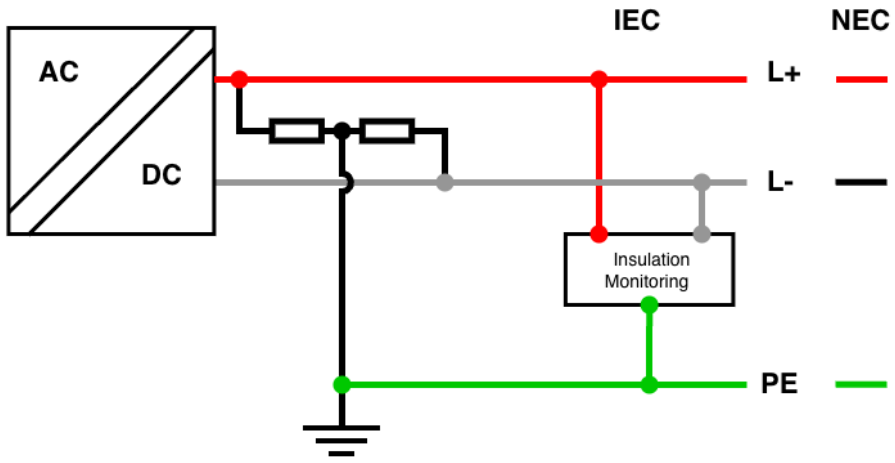


Figure 6.11. IT System ground reference with symmetrization and insulation monitoring

6.2.3.1. Direct Touch Protection (with IMD)

IT systems require a monitoring device to detect when a first fault to ground happens. An Insulation Monitoring Device (IMD), is a device commonly used to do this function. Fault locator devices are not mandatory but can be useful to locate the faulty feeder or branch. While it is

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

possible to continue operation under a first fault, it is not recommended to keep this situation for a long period of time, especially in the case of single pole protection, as it increases the risk of a second fault. Additionally, in an IT grounded system, only one IMD can be actively used at a time because multiple active devices will interfere with each other's measuring signals, ultimately leading to incorrect measurement results. Because there is only one IMD allowed to be active at a time, the single device sees the cumulative leakage current of the entire system (or facility). IMDs can also be used to measure the cumulative value of capacitance (line-to-earth) in the system.

6.2.3.2. Y-Capacitor Limitations

Y-capacitors provide electromagnetic interference (EMI) filtering and are manufactured to the IEC 60384-14 safety capacitor specification, ensuring that they fail open to protect users from electric shock. In smaller or low-power systems, the leakage current and stored energy are typically low enough that the equipment remains safe to touch. However, in larger systems with higher capacitance values, Y-capacitors can store significant charge, and during discharge this may create a hazardous touch voltage, posing a potential safety risk if not properly limited or controlled. The power levels discussed for direct current distribution in data centers exceed the levels that are non-problematic.

An existing example of this safety constraint can be seen in DC Electric Vehicle (EV) charging applications, which operate as ungrounded (IT) systems. To protect EV charging equipment users from electric shock, international safety standards limit the maximum allowable Y-capacitance. For instance, ISO 5474-3:2024 limits the Y-capacitance on the vehicle side to a maximum of 4 μF and the Combined Charging System (CCS) ≤ 400kW, IEC 61851-23 mandates that the maximum total parallel Y-capacitance on the charging station side must not exceed 1 μF (distributed as ≤ 500 nF across each DC rail to ground). Together they limit the total Y-capacitance in the charging circuit to 5μF. The CCS standard also requires the voltage between any live rail to ground to be not greater than 704V for a 1000V system . These voltage and capacitance limitations ensure the charging circuit remains below the b and c1 limits in IEC60479-1 Fig 22 and IEC60479-2 Fig 23.

As the electric vehicle charging industry scales to more than 1MW with the Megawatt Charging Systems (MCS), the capacitance required for proper electromagnetic interference (EMI) filtering naturally increases. This inherently pushes against these established safety thresholds, requiring more strict compliance requirements and making direct touch safety in the IT earthed

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

cable between charger and car a complex engineering challenge. The upcoming IEC 61851-23-3 increases Y-capacitance system limit from 5uF to 30uF for pole-to-pole voltage under 1078V with symmetrization (i.e. 539V line-to-earth). This results with the assumed body resistance of 575 Ohm to a current of ca. 500mA after 10ms, which is at the limit of the c1 zone in IEC 60479-2. The same IEC 61851-23-3 reduces the maximum allowed Y-capacitance from 30uF to ~11.4uF when the system voltage increases to max 1250V, and mandates active balancing of the voltages between live rails and ground. At 1500V operation in a perfectly symmetric system (750V line-to-earth) with a human body modeled by a single 575Ω resistor (as in IEC 61851-23-3 ED1), a 2uF total Y-Cap limitation in the IT circuit will keep the touch energy at the “c1” limit.

Therefore, it is good to assume that a large-scale interconnected direct current data center will require EMI filters with Y-capacitors of a total size that will make it impractical to make an IT earthed electrical system touch safe.

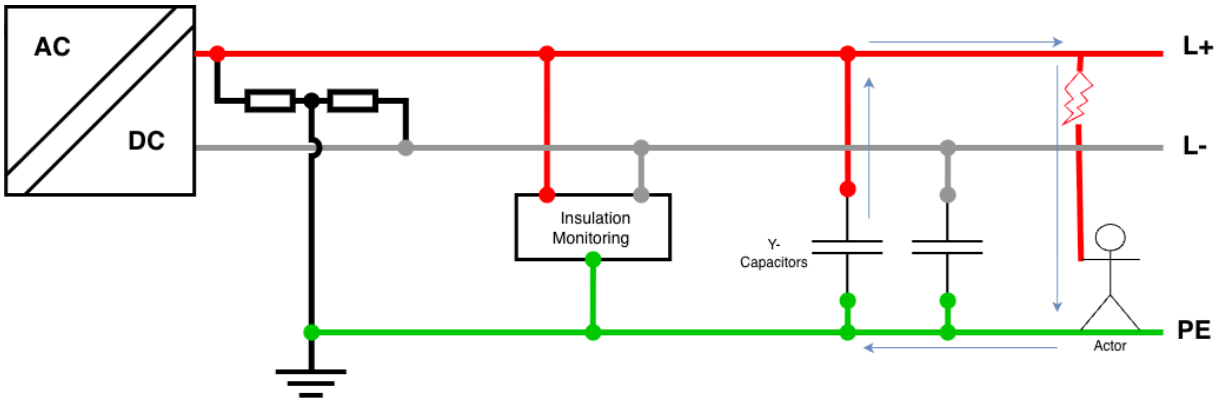


Figure 6.12. IT earthed system where system Y-capacitance discharges into human touch fault and resulting fault current path

6.2.4. Corrosion

Electrochemical corrosion is caused by currents at the point where they pass from a metal into an electrolytic medium, e.g., concrete or soil. Electro-chemically, this process is known as a partial anodic reaction. For DC current to cause corrosion, a material-specific minimum potential difference at that metal-electrolyte interface must be exceeded. EN 50162 specifies a threshold of 200 mV for steel-reinforced concrete [DIN EN 50162:2005-05 “Protection against corrosion by stray current from direct current systems”.]. Using Faraday’s law, it can be calculated that exposing an electrolyte to a current of 1 A for one year will cause the loss of 9.13 kg of an iron electrode (often used for providing the earth reference).

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

This type of corrosion can occur in DC installations when currents in the grounding system (PE) unintentionally – but sometimes unavoidably – leave the intended metallic PE conductor. One example is fastening (grounded) metallic housing to the foundation and the mounting screws coming into contact with the reinforcement steel of the foundation. An effective mitigation method against electrochemical corrosion is to ensure that the potential drop between any metallic housing and the central earthing point stays below the 200 mV threshold mentioned above.

In systems with an active wire at ground potential, such as unipolar systems or bipolar systems with a midpoint conductor (see chapter 6.1) an accidental connection of this midpoint wire to ground poses this corrosion risk. Means to detect this accidental connection require extra effort, such as residual current monitoring (RCM).

6.3. Combinations of Polarities, System Earthing and Voltages

Low voltage DC grids in data centers will use combinations of the various concepts introduced earlier in this chapter. In the following table, many relevant aspects for system choices and their impact on parameters and device compatibility are compared. Some of the aspects can be representative examples for other parameters that, although not shown, change similarly amongst the columns. In general, a power system or DC installation perspective is adopted for the parameters on the table rather than a power supply specification when considering the typical LVDC distribution in the data center. The Nominal Voltage is chosen from the 350V, 700V and 1400V, in order to scale consistently and indicate compatibility between scaling levels. This is also inline with international standardization for LVDC in IEC 60038 (draft), Current/OS, ODCA and EMerge Alliance. 380V, 400V, 750V, 800V and 1500V that have so far been used in the DC data center context fall into these classes. These equivalents are marked under the nominal table title row for ease of reference. The slash rated 350/700V and 700/1400V systems imply the presence of an M (or N in NEC) conductor that is at least partially distributed.

The grey columns are deemed less likely to be used in data centers. The columns with the white background are deemed likely to be used by at least one participant in the work stream.

A more detailed explanation and interpretation for the table is provided afterwards.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

Table 6.6. Combinations of voltages, polarities and system earthing and their impact on selected characteristics and device compatibility.

Nominal Voltage (equivalent) (V)	**350/700** (400/800)		**700** (800)				**700/1400** (750/1500)		**1400** (1500)			
Polarities	**Bipolar**				**Unipolar**		**Bipolar**				**Unipolar**	
Earthing	**TN-S**	**IT**	**TN-S**	**IT**	**TN-S**	**IT**	**TN-S**	**IT**	**TN-S**	**IT**	**TN-S**	**IT**
Design steady state nominal line-to-earth voltage (V)	350	700	350	700	700	700	700	1400	700	1400	1400	1400
Impulse Withstand Voltage, OVC I (kV)[a]	1.5	2.5	1.5	2.5	2.5	2.5	2.5	4	2.5	4	4	4
Impulse Withstand Voltage, OVC II (kV)[a]	2.5	4	2.5	4	4	4	4	6	4	6	6	6
Impulse Withstand Voltage, OVC III (kV)[a]	4	6	4	6	6	6	6	8	6	8	8	8
Impulse Withstand Voltage, OVC IV (kV)[a]	6	8	6	8	8	8	8	12	8	12	12	12
Clearance to PE, OVC II (mm)[b]	1.5	3	1.5	3	3	3	3	5.5	3	5.5	5.5	5.5
Clearance to PE, OVC III (mm)[b]	3	5.5	3	5.5	5.5	5.5	5.5	8	5.5	8	8	8
Clearance to PE, OVC IV (mm)[b]	5.5	8	5.5	8	8	8	8	14	8	14	14	14
Creepage to PE, PD2, for Mat Grp 1 (mm)[c]	2.5	5	2.5	5	5	5	5	8	5	8	8	8
Creepage to PE, PD2, Mat Grp 3 (mm)[c]	5	10	5	10	10	10	10	16	10	16	16	16
Creepage to PE, PD3, Mat Grp 3 (mm)[c]	8	16	8	16	16	16	16	25	16	25	25	25
Creepage for L-M, PD2, Mat Grp 1 (mm)[c]	2.5	2.5	-	-	5	5	5	5	-	-	8	8
Creepage for L-M, PD2, Mat Grp 3 (mm)[c]	5	5	-	-	10	10	10	10	-	-	16	16
Creepage for L-M, PD3, Mat Grp 3 (mm)[c]	8	8	-	-	16	16	16	16	-	-	25	25
Creepage for L-L, PD2, Mat Grp 1 (mm)[c]	5	5	5	5	-	-	8	8	8	8	-	-

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

Nominal Voltage (equivalent) (V)	**350/700** (400/800)		**700** (800)				**700/1400** (750/1500)		**1400** (1500)			
Polarities	**Bipolar**				**Unipolar**		**Bipolar**				**Unipolar**	
Earthing	**TN-S**	**IT**	**TN-S**	**IT**	**TN-S**	**IT**	**TN-S**	**IT**	**TN-S**	**IT**	**TN-S**	**IT**
Creepage for L-L, PD2, Mat Grp 3 (mm)[c]	10	10	10	10	-	-	16	16	16	16	-	-
Creepage for L-L, PD3, Mat Grp 3 (mm)[c]	16	16	16	16	-	-	25	25	25	25	-	-
SST Isolation *Recommendation*[d]	BI	RI	BI	RI	BI	RI	BI	RI	BI	RI	BI	RI
AC/DC Isolation Required if Connecting to LV AC TN-S	no	yes	no	yes	yes	yes	no	yes	no	yes	yes	yes
350V (or 400V) Device Compatibility	(x)[e]	(x)[e,f]	-	-	-	-	-	-	-	-	-	-
700V (or 800V) Device Compatibility	x	(x)[e]	x	(x)[e]	(x)[e]	(x)[e]	(x)[e]	(x)[e,f]	-	-	-	-
1400V (or 1500V) Device Compatibility	-	-	-	-	-	-	x	(x)[e]	x	(x)[e]	(x)[e]	(x)[e]
Active Wires	2-3	2-3	2	2	2	2	2-3	2-3	2	2	2	2
Relative Active Wire Material	2-3	2-3	2	2	2	2	1-1.5	1-1.5	1	1	1	1
Relative Cable Losses for same Balanced Power	1	1	1	1	1	1	0.5	0.5	0.5	0.5	0.5	0.5
Relative Current for same Balanced Power	1	1	1	1	1	1	0.5	0.5	0.5	0.5	0.5	0.5
Min. # of Overcurrent Protected Poles	2	3[h]	2	2[h]	1	2[h]	2	3[h]	2	2[h]	1	2[h]
Relative Voltage Semi-conductor Circuit Breaker[i]	0.5	1	1	1	1	1	1	2	2	2	2	2
Semiconductor Circuit Breaker FoM[j]	1	3	2	2	1	2	1	3	2	2	1	2

a) Impulse Withstand Voltage(Uw, Uimp): IEC 60364-4-44, currently under revision and a new table with DC values is proposed. The DC voltages are proposed as 1.5 times the AC rms voltage.
b) Clearance: IEC 60664-1, Table F.2, Case A inhomogeneous field. Values are for Basic Insulation (BI). OVC IV included for SST applications. Tables that define clearance distances in IEC 62368-1 also make reference to IEC 60664-1. Other relevant product standards (incl. NEC and UL) and their relation to the values in the table is shown below in section 6.3.1.1.
c) Creepage: IEC 60664-1, Table F.5, values for pollution degree 2 and 3 are shown (Note: use of ribs can reduce PD3 requirement), Material Group (Mat Grp) 1 for PCB design, Material Group 3

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

for Assembly. Values based on Vrms at 500V, 1000V, 1600V for 350V, 700V and 1400V, respectively. Values are for Basic Insulation (BI). Tables that define creepage distances in IEC 62368-1 also make reference to IEC 60664-1. Other relevant product standards (incl. NEC and UL) and their relation to the values in the table is shown below in section 6.3.1.1.

d) Recommendations for insulation coordination based on IEC 62477-1,-2, for HF TX between primary and secondary. (RI) Reinforced Insulation, (BI) Basic Insulation. In general, if it is accessible to humans it must be RI, if a human cannot access - it could be BI and if it can be tripped fast (RCD) - it could be BI. For a TN-S DC grid and there is SSCB, you could define this as BI with or without human accessibility. Topic is still under discussion and SST application based standards are still under development for this topic. See section 6.3.1.2 for more information.
e) Requires 1x the rated voltage and corresponding electrical spacing of the device from earth. See section below section 6.3.2 Combination Table Analysis: Device Compatibility for detailed description and legend, see Table 7.2for breakdown of x and (x).
f) Requires 2x the rated voltage and corresponding electrical spacing of the device from earth. See Table 11 for breakdown of x and (x).
g) With (partially) distributed M. For some characteristics this results in a range from having M distributed everywhere, to only having it in some circuits, and to no distribution of M but just usage in protection cabinets.
h) NEC 240.15(A) requires all ungrounded poles (regardless of TN-S or IT system type) to be protected, this is similar for IT requirements in references such as the IEC 60364-4-43 and NEN 1010. It is considered that IT earthed systems are designed for continued operation (for a limited time) after the first ground fault.
i) Semiconductor voltage example reference: 0.5=600V, 1=1200V, 2=2400V). Semiconductor selection (in a protection device) may vary by design. See section below titled section 6.3.4 Combination Table Analysis: Impact on Protection Devices for more information.
j) This Figure of Merit (FoM) quantifies the total semiconductor "effort" or cost required to protect a system. It is calculated by the multiplication of relative Semiconductor Circuit Breaker Voltage, Relative Current for same Power and Number of Protected Poles. For the current, the balanced case is considered in Bipolar systems with M. It is assumed that unbalance typically only occurs under lower load conditions.

6.3.1. Combination Table Analysis: Overvoltage Category (OVC), Impulse Withstand Voltage (Uw, Uimp), Clearance, and Creepage

Overvoltage Category (OVC), isolation type, and the maximum Working Voltage determine the Impulse Withstand Voltage (Uw, Uimp) which is the transient surge voltage, which will impact Clearance. Various clearance values as well as impulse withstand voltages are shown for various OVCs in the combinations table. Overvoltage categories are defined in standards such as IEC 60664-1 (ch. 4.3.2), 60364-1, and the respective product standards (e.g. IEC 60947, IEC 62477). The following table shows when different pollution degrees (PD) may be chosen in data center applications (depending on the location of the equipment within the datacenter).

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

Table 6.7. Pollution Degree (PD) Applications (IEC 60664, IEC 62368, UL 840, UL 508A)

PD	Applications (likely data center applications in grey)
1	Within fully sealed (encapsulated) equipment, components, and subassemblies. No pollution, or only dry, non-conductive pollution.
2	Default for IT equipment racks, servers, and power distribution units. Heaters within enclosures help to achieve PD2. Normally, only non-conductive pollution is present. Some condensation possible.
3	Ancillary facility zones (mechanical rooms, loading bays in DCF etc..). Conductive pollution or dry, non-conductive pollution that becomes conductive due to condensation.

The maximum working voltage, presence of potting or conformal coatings, isolation type, and pollution degree (describes how conductive contaminants in the environment affect insulation) impact the Creepage. Additionally, for the given applications (PCB design and assembly), the Material Group, a classification based on comparative tracking index (CTI), is defined. For PCB design, this may qualify as Material Group 1 assuming high CTI laminates or conformal coatings are used on the PCBs. It should be noted that certification bodies may require additional testing to claim this designation which may result in additional cost. As a result, Material Group 3 should be referenced for standard PCB FR4 (no coatings). A conservative approach is assumed for assemblies in DCF applications, providing a higher safety margin over the lifetime of the system. For this reason, Material Group 3 is taken as the spec for the assembly (such as busbars, terminal blocks, switchgear) level. Lastly, for industrial applications, an assembly could be taken to have lower Material Group 1 classification if proper material quality, laminates, insulation, and overall CTI requirements are achieved and validated by testing.

It is important to note that the creepage spacing cannot be less than the clearance spacing requirement for a given application. Therefore, the minimum creepage for a given system must be greater than or equal to the clearance.

6.3.1.1. Additional Discussion: Relationship between product standards and electrical spacing for NEC, UL, and IEC

In general, the majority of the IEC product and system standards refer to the IEC 60664-1 for electrical spacing requirements. Two relevant product standard examples are the IEC 62477 series (power electronics/converters) and the IEC 60947 series (including Part 10 for switchgear and semiconductor circuit breakers). Moreover, the UL 840 (Insulation Coordination Including Clearances and Creepage Distances for Electrical Equipment) is based on and harmonized with the electrical spacing criteria and values to those provided in the IEC 60664-1. However, other

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

product standards, particularly non-harmonized UL/NEC standards, can establish different spacing criteria which differ significantly from the UL 840. For example, UL 489i imposes stricter clearance and creepage distance requirements for SSCBs (e.g. 12.7 mm and 25.4 mm, respectively at terminals over 300V). Unlike IEC standards, UL 489i does not distinguish between pollution degrees (PDs) or material groups for creepage distances (except for internal spacings of PWB) as the SSCBs covered are for general applications throughout the electrical distribution system. Instead, UL 489i more closely resembles the more conservative spacing requirements for electrical distribution applications such as those referred to in IEC 61439 and UL 857. Additionally, the NEC (NFPA 70) Article 408 (see table 408.56) is aligned with the higher electrical spacing requirements in product standards such as UL 891 (switchboards), UL 67 (panelboards), and UL 1558 (low voltage power circuit breaker switchgear) to dictate fixed physical gaps strictly based on voltage levels, with no provisions for spacing reductions based on Overvoltage Category (OVC), Pollution Degree (PD) or Isolation Type (e.g. Basic, Functional). Some factors leading to the more conservative spacings from NEC/UL include: limited protection due to service equipment use, standards covering general application including areas with higher pollution degrees and possibility of contamination from external sources, termination of field conductors, long term aging of equipment, ionized gas release from interruption, electromagnetic displacement and mechanical deformation effects caused by high currents in applications with high incident fault energy, or when using fuses, mechanical circuit breakers, or protection devices with slower reaction times.

6.3.1.2. Additional Discussion: Isolation requirements in system design

Isolation Type (Basic, Functional, Supplementary, Reinforced) determines how much protection is needed and affects all related insulation test and design voltages (such as impulse withstand voltage). For example, reinforced insulation (RI) typically doubles clearance and creepage required under normal basic insulation (BI) conditions. For this reason, the table provides Basic Insulation (BI) values; if system design requires Reinforced Insulation (RI), these values can generally be doubled. Furthermore, BI is typically sufficient when the system chassis is properly earthed. This is highly advantageous for system architects, as the reduced clearance and creepage requirements of BI optimize board/cabinet space and increase spatial efficiency. It is also important to note that insulation requirements often vary throughout a single product or system. For instance, in a system utilizing an isolation transformer, the ungrounded side may require RI, while the properly grounded side may only require BI. Additionally, if conformal

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

coatings or overvoltage protection devices are used in the system, there may be locations where functional isolation (FI) is appropriate. All of this can be determined in application specific standards in IEC where isolation requirements are defined based on various safety, grounding, and isolation barrier considerations.

6.3.1.3. Additional Discussion: Voltage with respect to earth for different systems (polarities)

One clear trend that can be seen is that for bipolar TN-S systems, the clearance and creepage with respect to earth can be half of the value compared to a bipolar IT system. This is also half of the spacing required for the unipolar case to achieve the same voltage. Therefore, systems and devices can be reduced in size when using bipolar TN-S if they are designed properly with respect to earth and the other lines.

6.3.1.4. Additional Discussion: Line-to-line clearance

Line to line (L-L) clearance is not explicitly defined in the table because the values of L-L clearance do not depend on the polarity and system earthing. One possible interpretation is that this L-L clearance can be based on minimum required clearance based on the working voltage e.g 800Vdc. The other interpretation is that a full non-direct impulse could still appear between the lines, thus you would need to define the impulse withstand voltage based on double the value of working voltage used for clearance for L-M. Effectively, in this case, there is roughly double the value of L-M clearance for L-L clearance.

6.3.2. Combination Table Analysis: Device Compatibility (350V or 400V device compatible, 700V or 800V device compatible, 1400V or 1500V compatible)

The purpose of these rows is to provide multiple perspectives on device compatibility in various system configurations for various device voltage ratings. These rows can be interpreted from two different perspectives; a product designer / manufacturer perspective and a system integrator, architect, or data center operator perspective.

- From a product designer / manufacturer perspective, engineering a device to meet the 700V compatibility range ensures it can be deployed across all indicated (see Table 6.6) system configurations. This provides designers with a clear roadmap of compatible

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

environments for their hardware. In addition, it shows that by simply increasing the internal electrical spacing, a device can meet the requirements for a broader group of systems, maximizing its market versatility.

- From a system integrator, architect, or data center operator perspective, the selection of a specific voltage, polarity, and earthing configuration establishes the technical boundaries for the facility. These choices define the group of compatible, voltage-rated products available for use, thereby providing a structured framework for long-term strategic planning and infrastructure design.

Table 6.8: Legend for the following rows in Table 6.6: 350V or 400V device compatible, 700V or 800V device compatible, 1400V or 1500V compatible. Referenced in Footnote (e) in Table 6.6.

Indicator	Description
x	Any product designed with the electrical spacing (clearance and creepage) for the indicated voltages will be a compatible device in the system for a given polarity and system earthing configuration
(x)	Special considerations have to be taken into account for cross-compatibility. Specifically, the full voltage to ground has to be considered in the design for spacing clearance and impulse voltages etc. For example, devices only designed for bipolar 350/700V might not be usable for unipolar 700V if not designed with sufficient (larger) clearances towards earth. The (x) label indicates additional caution should be taken to properly design products such that they can be used over a wider range of possible distribution system configurations.
-	Voltage is not available in the system for device to be used

6.3.3. Combination Table Analysis: Wiring (Active Wires, Active Wire Material, Relative Cable Losses, Relative Current for same balanced Power)

These rows show that for a given power, how the active wires, relative active wire material, cable losses and current rating compare under various configurations. These rows show that transitioning from 350/700V to 700/1400V with 0.5 scaling of active wire material still allows for the relative cable losses and relative current rating to both be halved, while maintaining a constant delivered power in all cases.

An additional consideration is also shown for active wires where a higher number of active wires, in some cases, can have positive efficiency impacts (e.g. 700V unipolar TN-S vs.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

700/1400V bipolar with M TN-S). In a bipolar setup, the third wire (known as the middle/midpoint conductor (M) or neutral conductor) only handles the leftover "imbalance" between the two poles; if the loads are equal, almost no current flows through the M wire. This drastically reduces overall losses for the same power, compared to the unipolar case.

6.3.4. Combination Table Analysis: Impact on Protection Devices

The minimum number of protected poles from overcurrents (incl. overload and short circuit) (Row label in Table 6.6: Min. # of Overcurrent Protected Poles) is a protection device related consideration, where a higher number signifies that additional overcurrent/overvoltage coordination may be required. It does not necessarily mean that more protection devices are required. As an example, semiconductor circuit breakers can be designed to connect on a single line (pole) or as standalone devices that connect over multiple lines. Moreover, bipolar systems with M might require more protected poles but, in unipolar systems, double the number of breakers would be required to connect the same number of circuits/loads if devices rated for two different voltage levels are required in a single system. Additional points on protection related analysis for Table 6.6 are shown below:

- The Relative Voltage Semiconductor Circuit Breaker row in Table 6.6 (SSCBs - solid-state circuit breakers) is a SSCB consideration (see footnote (i)). It specifies the relative voltage rating for SiC and GaN semiconductors required in various configurations. It shows that some semiconductor breaker designs require double the voltage for a given application. Typically, the lower the voltage rating, the lower the cost for a given semiconductor.
    - The Semiconductor Circuit Breaker Figure of Merit row is explained in footnote (j) and quantifies the total semiconductor "effort" or cost required to protect a system. It should be noted that this FoM is specifically impactful for SSCBs because a conventional electromechanical breaker does not have high cost scaling based on voltage for the LVDC level of voltage rating devices. Therefore, the cost for protection devices for electromechanical breakers and fuses only scales with the number of devices, while an SSCB also scales with voltage rating.
- The current rating (Relative Current for same Power row) from a circuit breaker perspective, means that if more current is required for the same power, larger current rated circuit breakers are required, which is a cost and system sizing impact.

More information on protection devices can be seen in Section 7, Protection and Safety.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

6.4. Conclusion

In this chapter, the polarities of DC grids were introduced. Unipolar and bipolar configurations, with and without (partially) distributed M conductor, have been described and characterized. In the system earthing section, firstly, electrical system boundaries created by isolation have been addressed. Then, the earthing options, such as TN-S and IT, and their implications have been discussed for these electrical systems. Notably, one of these implications is a misconception surrounding IT grounding and the assumption that it provides an inherently touch-safe environment for employees. This chapter concludes that actual touch safety is heavily dependent on the chosen system earthing scheme, voltage polarity, and facility scale.In standard ungrounded IT systems, the total system Y-capacitance must be strictly limited to prevent the cumulative facility from discharging a lethal shock through a human. Ultimately, this limits system sizing and signifies that IT grounding does not work for large-scale systems. This conclusion is supported by Y-capacitance limitations from leading EV charging standards. In addition, this scaling limitation can be severely worsened by system voltage. For example, an earth fault in a unipolar 800V IT system has a significantly higher voltage-to-earth compared to a 350/700V bipolar system. This higher fault voltage to earth proportionally increases both the discharge current and the total energy stored within the system's Y-capacitance. Bypassing these capacitance limits requires a grounded system (e.g. TN-S) and localized earth fault protection, like selective branch-level RCDs, to stop the facility's cumulative Y-capacitance from discharging during an earth fault. Because the multiple earth references required for selective RCD protection cause severe DC corrosion under standard hard-grounding methods, architectures can utilize System Referencing Devices (SRDs). SRDs block continuous DC but conduct high-frequency transients, successfully enabling localized protection, and safe scaling for next-generation DC grids. Lastly, the various combinations of polarities, system earthing and nominal system voltages have been compared and analyzed based on the impacts on various components.

The result of the analysis is that some combinations of these system options are more likely than others to be implemented in data centers. It should be noted that this could also change over time as power increases and new devices and components become available. In addition, the analysis shows that both data center architects and product manufacturers should account for synergies between possible data center architectures and the interoperability / cross-compatibility of devices being developed for future data center applications.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

The authors identified some options as likely to be used in data centers and others as less likely. In the following they are listed with some (non-exhaustive) considerations.

6.4.1. Likely to be used Combinations

The following options are considered likely to be used in data centers (not listed in any specific order). All likely combinations can supply 700V and 800V devices (if designed for full voltage to earth).

- **350/700V Bipolar TN-S with (partially) distributed M (400/800V)**
  - Here also 350V and 400V devices can be supported.
  - Has reduced voltage to earth compared to unipolar 700V for enhanced safety and can also connect without isolation to LVAC.
  - This system is described in the OCP Mt Diablo Specification.
- **700V Bipolar TN-S without distributed M (800V)**
  - Has reduced voltage to earth compared to unipolar for enhanced safety and can also connect without isolation to LVAC.
  - Only two active wires needed, no M conductor.
- **700V Bipolar IT without distributed M (800V)**
  - Reduced and more balanced leakage currents compared to the Unipolar IT system.
  - Option in the NVIDIA 800 VDC Architecture design document.
- **700V Unipolar TN-S (800V)**
  - Only one active pole has to be protected
  - Higher clearances required than 350/700V Bipolar TN-S.
  - Option in the NVIDIA 800 VDC Architecture design document.
- **700V Unipolar IT (800V)**
  - Bipolar IT without distributed M is very similar and reduces leakage currents.
  - This system also has higher Y-capacitance discharge currents.
- **700/1400V Bipolar TN-S with (partially) distributed M (750/1500V)**
  - In this case, the full LVDC range up to 1500V is used.
  - Both 700V or 800V and 1400V or 1500V devices could be supported.
  - The voltage to earth is the same as in unipolar 700V so the same devices could be used in both configurations.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

- Sources will need to ensure that steady state voltage stays below the 1500V low voltage DC limit, e.g. by limiting voltage to 750V per pole.
- 800V devices can only be used if they can operate under 750V. This means 800V with only +-5% operation range can NOT be used.

6.4.2. Less likely to be used Combinations

The following options are considered less likely to be used in data centers mostly because of better alternatives that are very similar mentioned above. Some reasons are provided per system (not exhaustive).

- **350/700V Bipolar IT with (partially) distributed M**
  - 350V devices would experience 700V to ground in case of first fault.
- **1400V Bipolar TN-S without distributed M**
  - Can not support 700V devices that are expected to be available.
- **700/1400V Bipolar IT with (partially) distributed M**
  - 700V devices would experience 1400V to ground in case of first fault.
- **1400V Bipolar IT without distributed M**
  - Can not support 700V devices that are expected to be available.
- **1400V Unipolar TN-S**
  - Can not support 700V devices that are expected to be available.
  - Bipolar TN-S offers lower clearances and creepage to enclosure.
  - In some countries LVDC has an additional limit of 900V line-to-earth.
- **1400V Unipolar IT**
  - Can not support 700V devices that are expected to be available.
  - Bipolar IT without distributed M is very similar and reduces leakage currents.
  - In some countries LVDC has an additional limit of 900V line-to-earth.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

7. Protection and Safety

7.1. Introduction

Protection of electrical installations -including data centers- from faults, such as short circuits, arcing, insulation or ground faults, is essential for ensuring safety and reliability. Without protective measures, faults can lead to severe consequences for people, equipment and property.

When designing protection for a data center many aspects must be taken into account, including:

- protecting against overcurrents, that could overheat conductors and damage electronic equipment;
- protecting humans from electric shock, electrocution and arc flashes, that could reach extreme temperature and cause severe burns or blindness;
- preventing electrical faults, including overcurrent and arc faults, that can evolve into fires;
- protecting sensitive electronic equipment from high-voltage disturbances such as transient voltage surges,
- ensuring continuity of service of the data center in case of a fault, by quickly isolating the faulted circuit and ensuring that operation in the rest of the installation remains unaffected (selectivity);
- meeting requirements of applicable codes and standards, including the National Electrical Code (NFPA 70) or wiring codes based on IEC 60364 standards.
- Protection against a critical failure of a semiconductor element imposing AC voltage on the DC system.

Unlike AC power systems and other DC systems, such as those in railways, DC power distribution systems for AI data centers are characterized by relatively low inductance and dominated by power-electronic sources and loads, e.g. AC/DC and DC/DC power converters or MV solid-state transformers (MV-SSTs). This changes the dynamic behavior and evolution of the faults and, consequently, the requirements on circuit protection. Due to the lower feeder inductances, resulting from the use of power electronics converters rather than transformers, and the contribution of DC bus capacitors discharging into the fault demand specific protection strategies that act within microseconds to milliseconds. Yet these challenges are well understood and have been effectively addressed through proven engineering solutions.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

In addition to traditional protection devices, fuses and electromechanical circuit breakers, new protection technologies - such as fast-acting, semiconductor (solid-state) circuit breakers and power converters with overcurrent blocking capability solid-state protection devices - have been introduced and have matured tremendously over the years.
This section provides a comprehensive yet practical overview of protection and safety in DC power distribution systems, based on best practices, real-world experience, and standards, aimed to prove that DC power systems to feed power hungry AI factories are safe and reliable and no longer a future aspiration, and to serve as a reference for protection design blueprints.

7.2. Protection against overcurrent

Overcurrent faults considered in the scope of this document are: pole-to-pole, pole-to-midpoint and pole-to-ground, short-circuit and high-impedance faults, including parallel arcs.
Protection from overcurrent in power systems is addressed in IEC 60364-4-43 and in Sections 110 and 240 of NFPA 70. As described below, power converters play an important role in fault protection, further investigations are needed if this is properly considered in current codes and standards.

7.2.1. Overcurrent faults in active DC systems

It is assumed that in DC power distribution systems for data centers, sources and loads are interfaced to the DC bus through power converters that regulate the voltage or the current and maintain the system in the required control band. These DC installations are characterized by capacitive sources connected at a short distance (hence with low inductance) to the loads, in the event of a short circuit the fault current increases very rapidly; an increase in the order of several tens of A/μs has been experienced (Fig. 7.1).
Currently, no standards or recognized practices exist providing methods for assessing the short circuit current of this kind of DC power systems; future TS 63713-1, under development by IEC TC73, is aimed at filling this gap.
Considering a simple system with only one converter, the short-circuit current exhibits an initial peak resulting from the discharge of the capacitors followed by the steady-state contribution from power conversion, which depends on the type of the converter (Fig. 7.1 and 7.2). Fault-limiting converters [Ca] (see Section 8.3) are able to limit the steady-state stage of the fault current to zero or low values in the range of the nominal current, not sufficient to sustain the voltage on the DC bus. If the fault is not cleared in time, the voltage decays to values lower

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

than the minimum operating voltage of the system; the time to reach this level depends on the capacitors and on the fault-loop inductance and can be lower than 100 µs. Fast-acting overcurrent protection devices can be used to limit the DC voltage drop and ensure unaffected parts of the system continue operation. An alternative protection strategy is to allow the DC bus voltage to drop below the minimum operating voltage for a certain duration, providing a sufficiently large fault current to trigger the protection devices and clear the fault, in the meantime the loads need to be supported by local energy storage; once the fault is cleared, the DC bus voltage can recover, and the loads will be supported by the source again. On the other hand, fault-feeding converters (see Section 8.3) can contribute to sustain the DC voltage, but they require an overcurrent protection device fast enough to avoid that excess current results in permanent internal damage or fires.

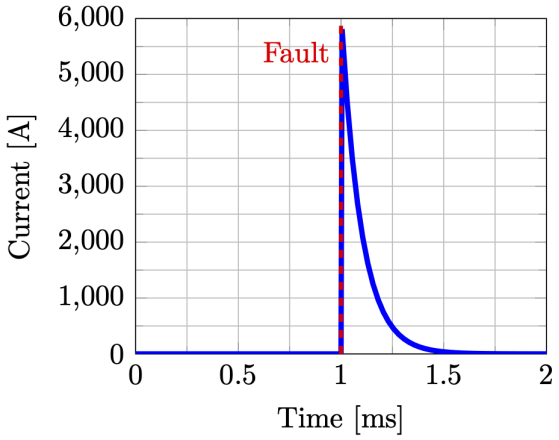


Figure 7.1 Fault current for a fault-limiting converter

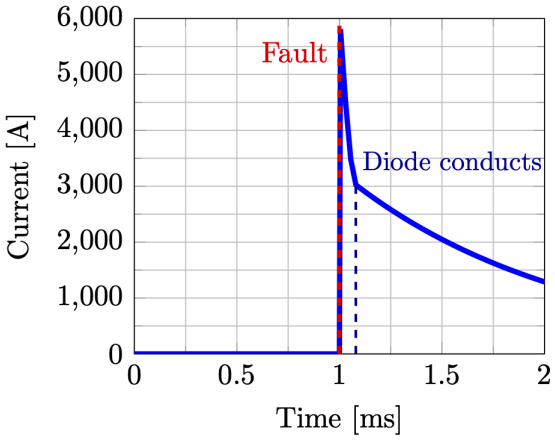


Figure 7.2 Fault current for a fault-feeding converter

Considering the capacitor discharge only, depending on the characteristics of the circuit, the fault current can quickly decay to zero (overdamped case, Fig. 7.3) or oscillate around zero (underdamped case, Fig. 7.4).

Aside from electrical and thermal effects, such high fault currents could result in electromagnetic force on current carrying conductors due to the Lorentz force. Conductors of the same polarity (i.e. with parallel flowing current) would be pushed together, whereas conductors of opposite polarities (with antiparallel flowing current) would be pushed away from each other, with intensity directly proportional to the squared current amplitude and conductor length and inversely with conductors' separation. Therefore, when high short-circuit currents of the order of tens or hundreds kA are expected, special care is to be taken by fixing busbars and conductors

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

in trays to avoid displacement and accidental disconnection, which eventually may lead to arcing. Ultrafast protection with SSCBs, would limit the fault current and mitigate the issue.

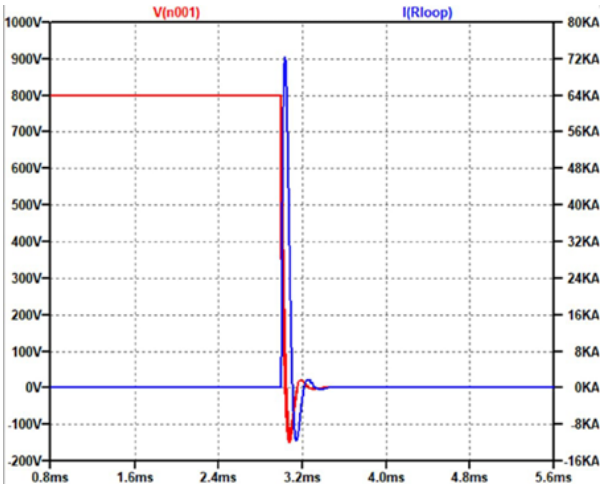


Figure 7.3 Fault current in an overdamped circuit

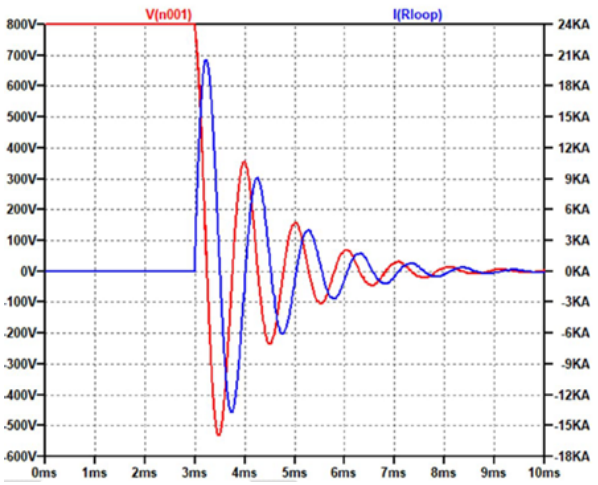


Figure 7.4 Fault current in an underdamped circuit

In more complex systems, multiple converters connected to the DC bus may contribute to the fault cur-rent. Also load converters, although not intended to inject current in the DC bus, can contribute to the fault current through the discharge of their input capacitors. To avoid this effect, power supplies in the IT-load side may have OR-ing diodes or switches, which prevent the input capacitors being discharged by the upstream faults, so that they can provide hold-up time in the case of upstream faults and voltage drop.

To demonstrate the qualitative behaviour described above, consider a representative 800 V DC reference design and following voltage and current waveforms in fault conditions obtained through time domain simulations. For clarity the analysis considers a single fault location: a bolted short circuit at one server, at the far end at the server.

Architecture. The architecture is a 4 MW AC service feeding an AC/DC converter onto an 800 V collector bus (Fig. 7.5). A main DC/DC converter (800/600 V) with a 600 V battery (BAT-1, 2 mΩ internal resistance) shares the collector bus. The bus feeds the main bus through the main breaker. The main bus distributes power to four feeders, each protected by a branch circuit breaker, containing an optional blocking diode and a load-side bus that carries one 100 kW server load and a secondary 600 V distributed battery behind a dedicated DC/DC converter. A

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

fault at one server is the instructive case: it exercises the faulted PSU bus local storage, the branch protection, and the interaction with the three healthy feeders.

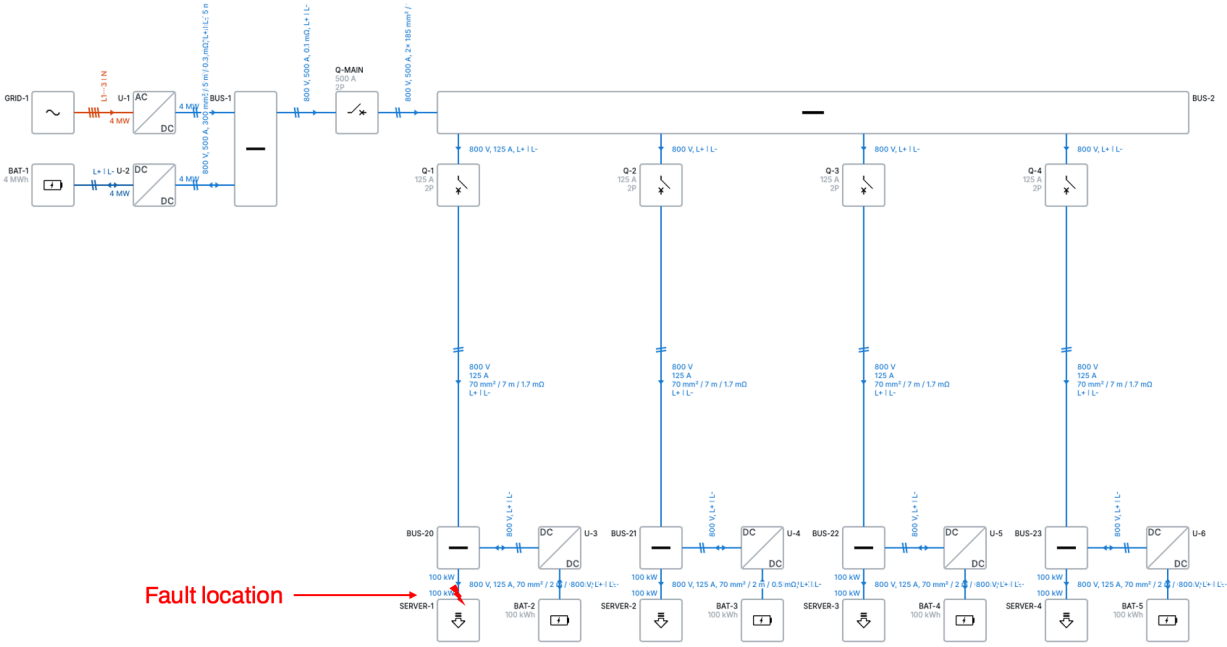


Fig. 7.5 Single-line diagram of the 800 V DC reference design. Full reference design topology and parameters available online [DCI].

Table 7.1 Model parameters

Parameter	Value
Fault	bolted, at t = 1 ms on fully energised system
AC-side equivalent (front end)	R = 1 mΩ, L = 20 uH per phase
Converter to source bus cable	0.3 mΩ, 1 uH
Source bus to main bus cable	100 m, 9.1 mΩ, 20 uH
PSU bus to server cable	2 m, 0.5 mΩ, 0.4 uH
AC/DC output DC link	70 mF, ESR 7 mΩ
Main DC/DC port DC link	90 mF, ESR 9 mΩ per port

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

Distributed DC/DC port DC link	10 mF, ESR 1 mΩ per port
Main-side battery (BAT-1)	600 V, 2 mΩ internal resistance
Rack battery (BAT-2 to BAT-5)	600 V, 20 mΩ internal resistance

Fault current and the role of fault-limiting converters. The server fault current rises extremely fast and reaches a high peak about 92 kA (Fig. 7.6), dominated by the discharge of the PSU bus and server DC link capacitances into the near-zero-impedance fault. The initial rate of rise is about 1 kA/μs, orders of magnitude faster than a comparable AC fault. The peak is set by the local capacitor discharge and is essentially the same whether the converters keep feeding the fault or limit their contribution; the steady-state level is what the converters set. A fault-feeding design keeps delivering and the fault settles near 85 kA. A fault-limiting design, in which every converter limits its fault contribution, drives the steady-state fault current to zero or near the nominal current. Reaching zero requires limiting the load-bus DC/DC as well as the source-bus converters: if only the source-bus AC/DC and DC/DC limit while the load-bus DC/DC keeps feeding, the faulted PSU bus's own battery still sustains the fault (600 V behind 20 mΩ gives about 30 kA), so the steady-state current falls to zero only once the load-bus converter is fault-limiting too. The conclusion is that fault-limiting converters, and fault-limiting behaviour in general, can control the steady-state fault current even when the DC bus has collapsed, holding it to zero or to the range of the nominal current. Fault-feeding converters cannot, and the large sustained current and its energy become a safety and equipment-rating concern that the protection must then handle.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

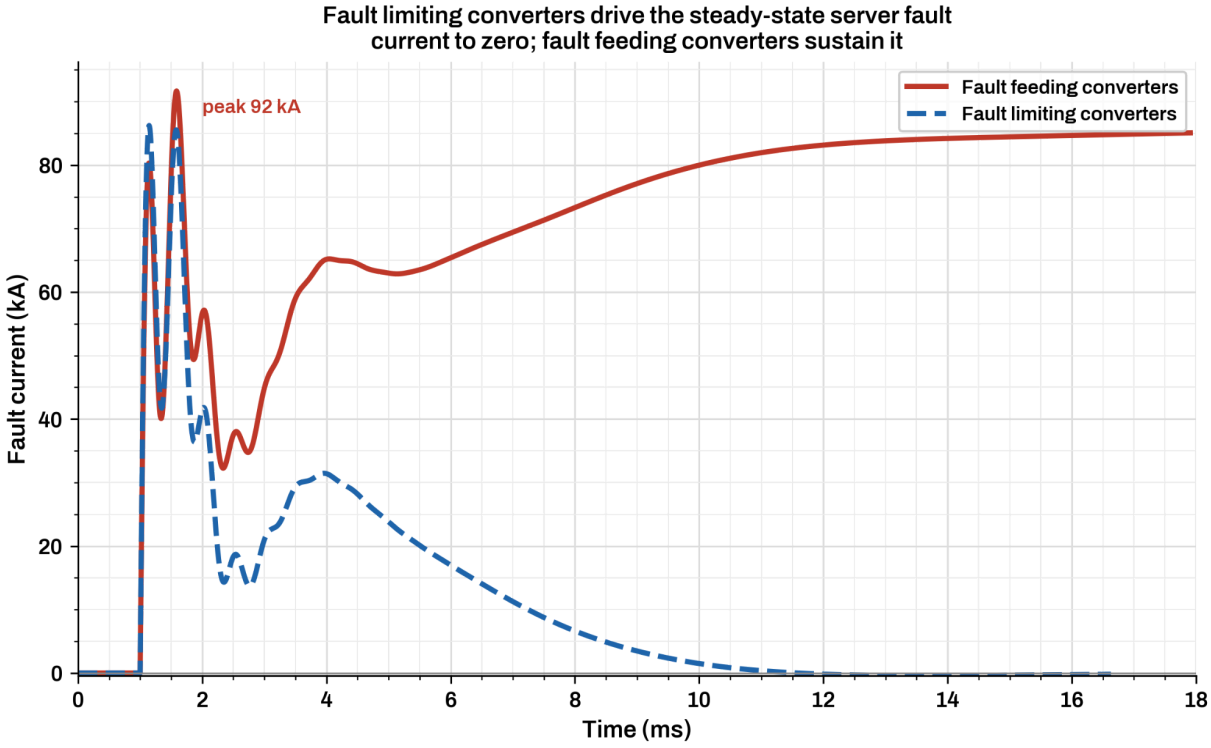


Fig. 7.6 Total fault current at the faulted server, fault-feeding versus fault-limiting converters. The capacitor-discharge peak (about 90 kA) is essentially the same; at steady state the fault-feeding design holds about 41 kA while the fault-limiting design, with every converter limiting including the load-bus DC/DC, reaches zero.

Fast voltage drop on the healthy branches. A server fault does not stay local: it pulls down the voltage on the neighbouring feeders, and it does so very fast. Fig. 7.7 shows both effects on one time axis, the fault current rising on the left axis and a healthy server's voltage falling on the right axis. The healthy pole voltage falls below 90% of nominal within about 0.26 ms of fault inception and below 50% within about 0.56 ms, reaching a minimum of about 60 V, while the fault current is still climbing toward its peak of about 92 kA. That collapse is faster than any conventional electromechanical DC breaker can clear (typically several ms), so either fast protection is required to keep the healthy part of the system operating through a fault elsewhere, or opt for a redundant power supply where losing the faulty branch does not jeopardize the system's availability.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

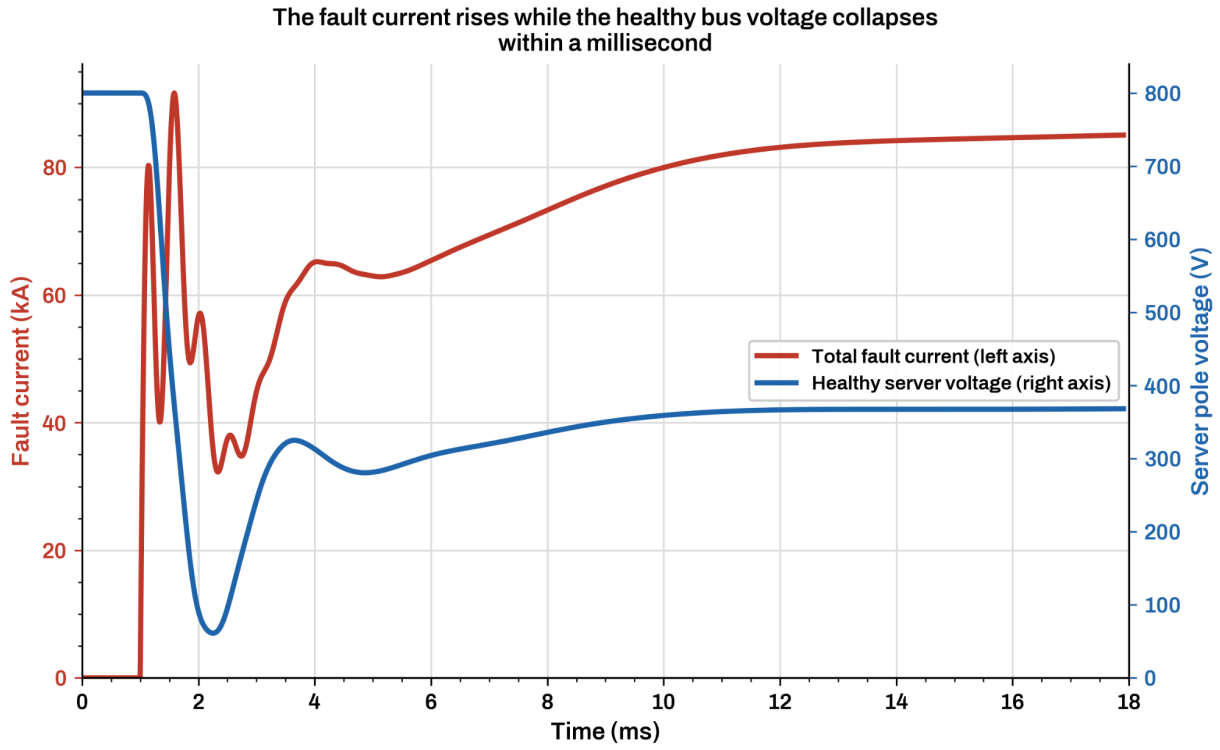


Fig. 7.7 Fault current (left axis) and the pole voltage at a healthy server (right axis) during the server fault, fault-feeding design with no reverse-path protection. The current rises as the healthy bus voltage collapses, both within a millisecond.

PSUs are sources of short-circuit current. A distinctive feature of DC systems is that the PSUs themselves can feed a fault. The total current at the faulted server is the sum of three contributions (Fig. 7.8): the source bus, delivering down through the main bus and the faulted feeder (about 27 kA at steady state); the neighbouring branches, back-feeding through their now-conducting feeders (about 33 kA from the three healthy load buses together); and the faulted load bus's own distributed storage (about 24 kA). At steady state these add to about 85 kA, the total fault current. Each server PSU presents a DC link capacitor that discharges into a nearby fault, and a distributed battery keeps pushing current behind it, so the PSUs, which normally only draw power, act as sources during the fault (Figure 7.8).

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

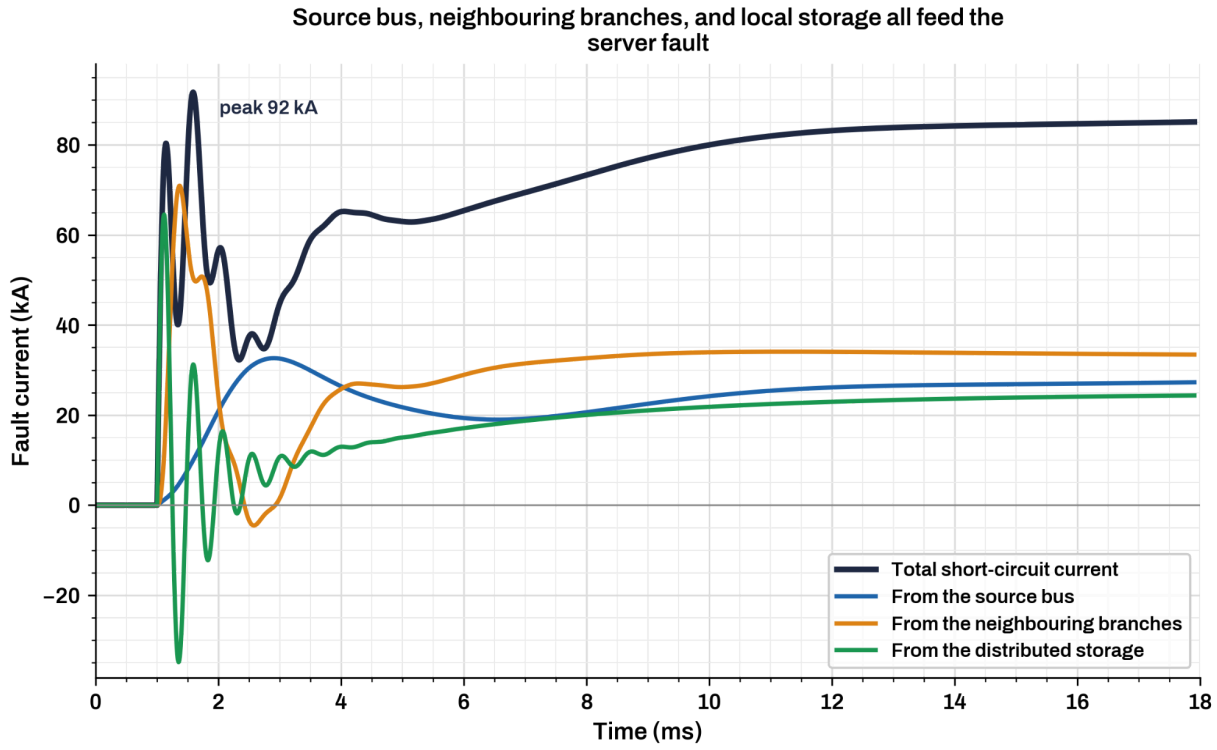


Fig. 7.8 The server fault current decomposed into its sources: the source bus, the neighbouring branches, and the faulted PSU bus's distributed storage.

Nuisance tripping. Because the neighbouring PSUs act as sources, each healthy branch carries a reverse current toward the fault (Fig. 7.9). It pulses to about 24 kA during the capacitor discharge and settles near 11 kA, fed by that load bus's storage and PSU input capacitors. A current-based protective device on a healthy branch sees this reverse current and can trip even though the fault is on a different branch, disconnecting a healthy PSU that should have stayed in service. This reverse current contribution is commonly underestimated in DC protection design: here the three healthy branches together back-feed about 33 kA into the fault, comparable to the source-bus contribution, so it is a first-order term that branch protection coordination must account for.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

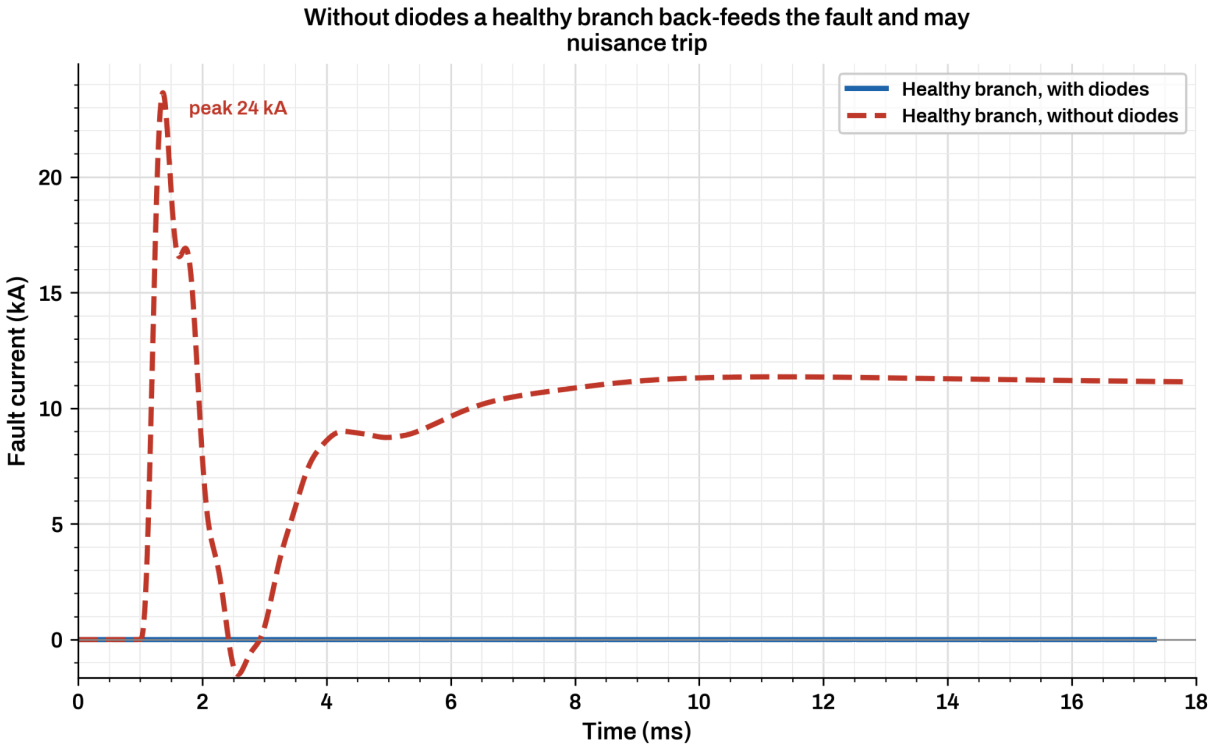


Fig. 7.9 Reverse current carried by a healthy branch during the server fault, the driver of nuisance tripping. The blocking diodes cut it to zero.

Diodes prevent the back-feed effect. A blocking diode in each feeder reverse-biases as soon as the main bus sags below its PSU bus, cutting the reverse path. The healthy branches then carry no reverse current, their breakers do not nuisance-trip, and their servers stay at 800 V for the whole fault (compare the with-diodes and without-diodes traces in Figs. 7.10 and 7.9). The diodes also keep the neighbouring PSU buses from adding to the current at the fault. This protection against neighbouring-feeder tripping and voltage collapse does not strictly require diodes: it can also be achieved, and has been verified, with semiconductor circuit breakers that detect the direction of current flow and apply a short time delay to currents flowing toward the DC distribution bus, which avoids the extra blocking diodes and their associated power loss at the cost of an active device [ODCA, VDE SPEC 90037].

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

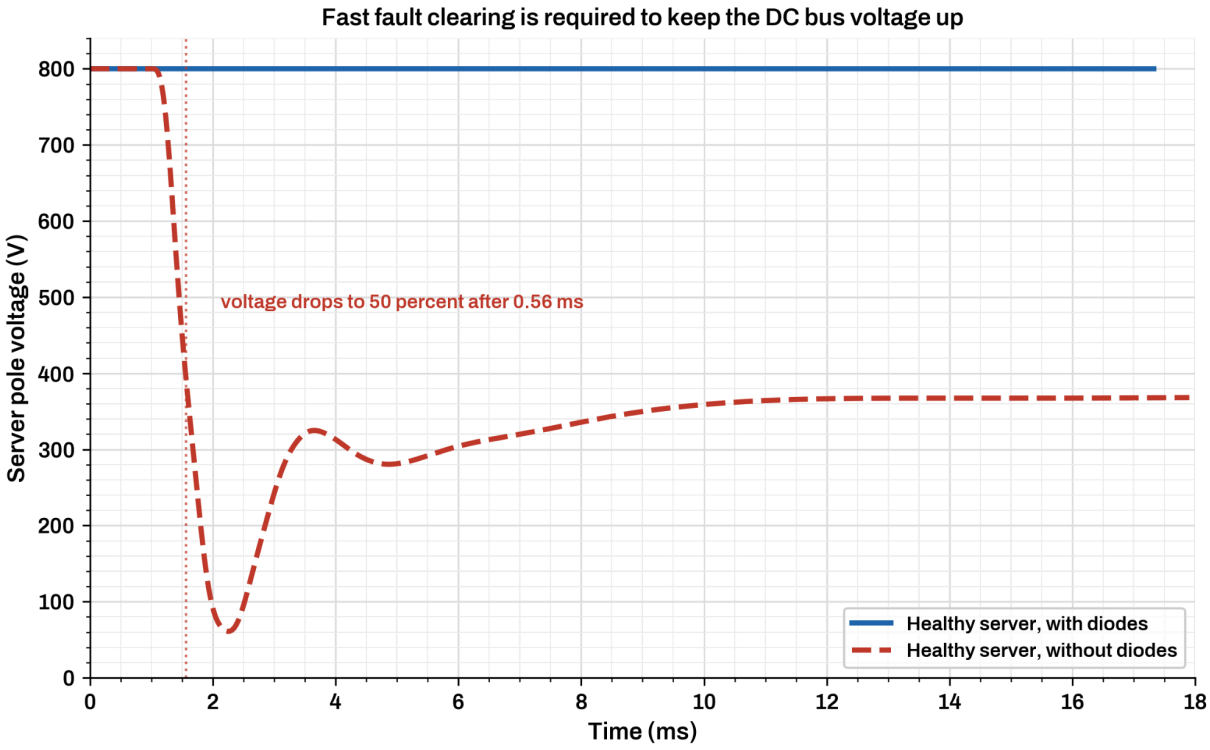


Fig. 7.10 Voltage at a healthy server during the server fault. With the blocking diodes it holds at 800 V; without them it collapses in under a millisecond.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

7.2.2. Selectivity

As mentioned above, protection is a key aspect not only to avoid injury to personnel and damage to the installation but also to achieve the high system availability required for data centers. In case of a fault, the healthy (fault-free) part of the system must, as far as possible, continue operation.

Protection selectivity (or discrimination) is a system's ability to have the circuit protection device closest to the fault interrupt and therefore isolate only the directly affected part of a system in the event of a fault, while ensuring continued operation of the fault-free components and sections as far as possible.

Selectivity is also the key to enable redundancy: failing parts of the system need to be isolated from the healthy system as soon as possible, to allow power to be re-established to the affected loads if possible.

Over the years, several methods have been developed to achieve selectivity in AC distribution systems, but Time-Current Selectivity remains the most widespread. This consists of setting protection devices to trip at different time delays for the same level of overcurrent, with the device closer to the load set to trip faster and the upstream device waiting longer. If the downstream breaker clears the fault in time, the upstream device does not trip. The high inertia of AC systems makes this relatively easy but can result in delays of up to 100-200 ms and, consequently, in case of arc flash, long arcing time and high energy release.

In a DC system, due to the fast fault development, in case of a short circuit the voltage can drop quickly. The fault should be cleared before the voltage drops below the minimum operating voltage of the converter and the system ceases to function, unless alternative strategies are implemented, as detailed below.

Additionally, with many energy storage devices connected in parallel, ensuring selectivity is more challenging than in a conventional linear configuration with a central power source, since in the event of a fault all energy storage devices will discharge into the fault location.

For selectivity, it must therefore be ensured that the protection devices in unaffected parts of the system will not trip when energy storage devices, such as capacitors, feed current into a fault via the DC grid (see 7.2.1), which can be achieved, for example, by detecting the direction of current flow in the DC breaker. Alternatively, back-feeding can be prevented by means of a diode.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

Some possible alternative methods that require further investigation include:

- use fast protection devices, e.g. fast semiconductor circuit breakers (see below) to avoid the voltage drops below the minimum voltage. This is the simplest solution, but semiconductor circuit breakers are just arriving in the market now and the product portfolio may not cover all needs in the shortest terms.
- Add energy storage (e.g. capacitors or supercapacitors (CBU), or battery back-up units (BBU) in rack to store enough energy to meet the Voltage hold-up time requirements with the breaking time of the existing protection devices. This solution is based on mature protection technology, but the amount of stored energy needed must be carefully evaluated to avoid safety and reliability issues.
- Redundant-feeders architecture, in which each rack is supplied by at least two independent DC buses, arranged in a way that a fault in a bus does not propagate to the others, which can continue feeding the rack. An example of this is in Figure 7.11, showing a 4-to-make-3 redundant-feeders architecture with four DC buses each supplying four PSUs in a rack with 50 V vertical bus. In normal operation, each PSU works at 75% of the rated power (a). In case of a fault on Bus A, regardless of whether there is coordination between the feeder protection device and the branch protection device (b) or not (c), the remaining PSUs, operating at 100%, can ensure service continuity of the racks. However, it's worth noting that coordination has a strong positive impact on reliability: without coordination, the system will fail in case of two faults occurring on two different buses, independently from their location; whereas with coordination the rack will fail only if the faults are in that same rack. Redundant-feeders architecture has less requirements in terms of protection, but it introduces some design constraints that can reduce flexibility.

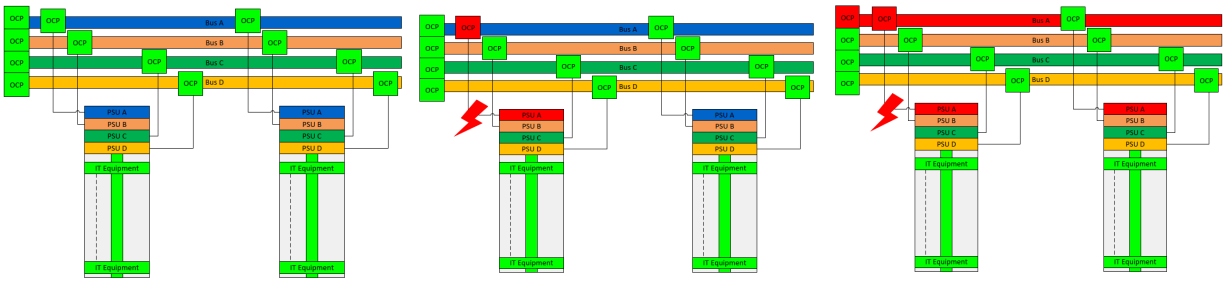


Figure. 7.11 4-to-make-3 redundant-feeders architecture. On the left (a), normal operation. In the center (b), fault with coordination of protection devices: the branch protection trips and the feeder protection does not. On the right (c), fault without coordination:both the feeder and the branch

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

In any case, if no intentional tripping delay is used, the arcing time in case of arc-flash may be limited to the time necessary for the protection devices to interrupt instantaneously, and consequently lower energy release can be expected.

7.2.3. Branch vs supplementary protection

NFPA 70 distinguishes between overcurrent protective devices based on many factors including design, ratings and application. Supplementary overcurrent protection devices are considered to provide limited overcurrent protection within equipment, and is only supplementary to protection provided by an overcurrent protective device in the distribution system. Overcurrent protective devices such as fuses and molded-case circuit breakers (also referred to as branch circuit overcurrent protective devices), are considered to provide protection to the distribution system, which includes branch, feeder and service conductors and equipment. Such distinction does not exist in IEC standards.

An example of supplementary protection protective devices are gR-class semiconductor fuses according to UL 248-13, intended for protection of power electronic equipment, which are not suitable for branch protection. Another example may be a hotswap controller located at the server shelf input. It protects the input connector from inrush current and arcing damage during hot-swap events, and it can also function as an electronic fuse to interrupt the current path during severe overcurrent or short-circuit conditions.

Requirements for branch circuit overcurrent protection devices are typically more stringent than for supplementary protection devices, e.g. they have higher minimum short circuit interrupting ratings, larger clearance and creepage distances, and more robust performance criteria. NFPA 70 specifies branch overcurrent protection devices to be readily accessible, whereas supplementary protection devices can be embedded in the equipment, e.g. in cabinets.

An example of a branch circuit overcurrent protection device is a circuit breaker located at a DC PDU output in the power rack or sidecar, or at the input of a DC bus. It shall be easily accessible, provide a clear visual indication of disconnection, and support lockout/tagout (LOTO) for maintenance activities.

7.2.4. Dependence of protection devices on grounding

Conductors that carry operational current and which have a potential difference to ground need to be protected against overcurrents as per IEC installation rules (IEC 60364-4-43). This is the case of L+ and L- conductors in two-wires and three-wires IT systems, as per IEC 60364-1,

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

clause 1.9.3.4, Figure. 28 and 29. Additionally, for three-wires IT systems, IEC 60364-4-43, clause 431.2.2 demands protection also for the midpoint conductor M, with the capability to disconnect *all* live conductors including M. Protection is not mandated instead for grounded conductors, irrespective they are polarity or mid-point conductors (see also IEC 60364-8-82, F.4.3.1), but only in case they are solidly grounded, as high-impedance grounding is considered as IT.

Similarly, NFPA 70 Section 240.15, requires overcurrent protection for ungrounded conductors; grounded conductors are not mentioned, therefore protection is not mandated for solidly grounded and, unlike in IEC, high-impedance grounded conductors. This difference between IEC and NEC rules needs further investigation.

A recap of the number of poles to be protected against overcurrent is shown in Table 6.6.

Additionally, in case the IT system is intended to continue operation after the first ground fault, the protection device shall be able to break the full system voltage on each single pole. For SSCBs according to IEC 60947-10, this means that the protected poles shall have individual-pole operational voltage U_{e1} not less than the maximum system voltage. This requirement is not currently addressed in UL 489I.

7.2.5. Overcurrent protection devices

Many different overcurrent protection devices have been developed for DC, ranging from traditional fuses and electromechanical circuit breakers, that operate by creating an arc that develops an arc voltage larger than the system voltage, to new equipment based on power semiconductor devices.

Every overcurrent protection device rated for the appropriate voltage and current levels is able to clear overcurrent faults, regardless of technology, thus ensuring safety.

According to NFPA 70 Section 110.10, overcurrent protection devices shall be coordinated with the available fault current. However, there is currently no standard and no recognized practice providing guidance about such coordination for the DC faults described above, in particular about how to consider the capacitor-discharge peak described above. A dedicated Task Group coordinated by ULS is currently investigating how to address this topic.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

7.2.5.1. Fuses

Fuses need a minimum energy to melt and interrupt the current. In datasheets, this is typically stated in terms of the Joule integral (I^2t) of the current required for melting the fuse: I^2t_{melt}. In case the I^2t of the fault (I^2t_{fault}) is lower than I^2t_{melt}, then the fuse will not melt. In case of fault limiting converters, where the available fault energy is stored in the capacitors, such capacitors shall be large enough to meet such criterion mentioned [Ray]. As I^2t_{fault} in DC can depend on the shape of the fault current, the trip time -and consequently the let-through current- of a fuse can vary significantly.

In short-circuit, fuses have an ranging from less than 1 ms, for fast fuses intended for semiconductor protection (IEC 60269-4, UL 248-13), to several ms.

Low-Voltage fuses are covered by IEC 60269 and UL 248 standards. UL standards for fuses suitable for branch circuit overcurrent protection are currently limited to 600V.

Presently, standards for fuses do not cover the extremely high rates of rise of the short-circuit current since the standards look at *time constant* to reach a (constant) prospective current. The conditions in capacitive distribution systems are not yet covered. Performances, including the interruption time, under such conditions need to be verified.

7.2.5.2. Electromechanical Circuit breakers

Electromechanical circuit breakers according to IEC 60947-2, UL 489 or UL 1066 are also suitable for DC. Most DC circuit breakers in the market have a thermal-magnetic trip unit and use different methods (e.g. with poles arranged in series) to build up a large enough arc voltage, forcing the fault current to zero. Breaking time in short circuit ranges from about 1 ms, for small miniature circuit breakers (MCBs), to a few tens of milliseconds, for power circuit breakers.

In case of active DC systems fed by fault limiting converters, the thermal-magnetic trip unit may not fit very well with the fault behavior, as the capacitor-discharge peak can be too narrow to trigger the magnetic protection, and the low fault energy can result in very long tripping time or no trip at all. Electronic trip units would fit much better for this application, but very few products are available in the market. Electronic trip units can also implement additional protections, e.g. Ground-fault protection of equipment (GFPE).

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

7.2.5.3. Solid state Circuit breakers(SSCBs), Semiconductor Hybrid Circuit breakers(SHCBs) and e-fuses

Solid state, or semiconductor, circuit breakers (SSCBs or SCCBs), according to the new UL 489I and IEC 60947-10, use semiconductor switching devices, e.g. IGBTs, IGCTs or MOSFETs, in the main current path, to interrupt current, instead of an arc. They include mechanical contacts to ensure galvanic isolation, as required for safe maintenance, and typically have an electronic trip unit.

In a short circuit condition, these devices have breaking time in the order of microseconds. By clearing the fault in the early phase of the capacitor discharge, they limit the voltage drop even in installations with limited capacitance. Due to such ultrafast interruption, SSCBs are also able to break very high prospective short-circuit currents. As such, the key performance indicator is not the breaking capacity (Ics / Icu in IEC 60947-2, or kAIC in UL 489), but the maximum rate of rise of the fault current (di/dt) that the circuit breaker can interrupt safely. Depending on the design, SSCBs can break currents up to several hundreds A/µs.

However, as technology for these devices is quite new, they are in general more expensive and have larger footprint compared to electromechanical circuit breakers. An additional drawback is that they have larger power loss and may require forced-air or even liquid cooling, for larger current ratings, which could make integration in electrical panels more complex.

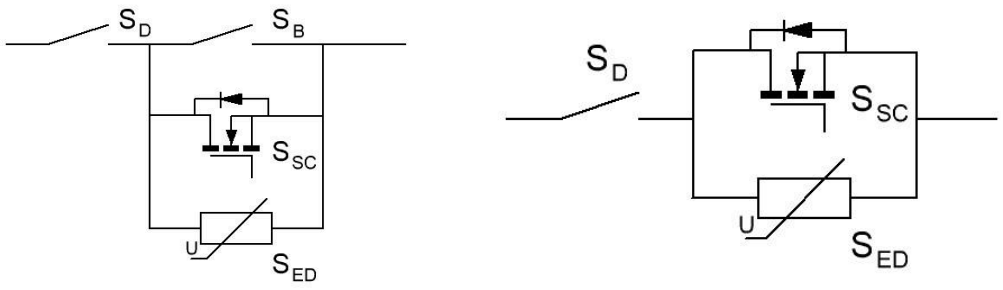


Figure. 7.12 concept design (left) for a semiconductor hybrid circuit breaker (SCHCB); (right) a semiconductor circuit-breaker (SCCB), also called solid state circuit-breaker concept (SSCB). Where: S_{SC} is the semiconductor switch, S_D is mechanical contact providing galvanic insulation, S_B is the by-pass contact, closed in on-condition to reduce power losses, and S_{ED} is the voltage clamping device, also dissipating the electromagnetic energy stored in the circuit.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

Semiconductor hybrid circuit breakers (SCHCBs), also covered by IEC 60947-10 and UL 489I, are similar to SSCBs but, in addition to the isolation contacts, they have mechanical by-pass contacts in parallel to the semiconductor switching device, intended for conducting the current in normal operation and therefore reducing power loss. SCHCBs have power losses slightly higher than electromechanical circuit breakers but, due to the inertia of the by-pass contacts (S_B in Figure. 7.12), they are much slower than SSCBs; their typical breaking time is in the order of a few hundreds μs, making them suitable for currents with a rate of rise up to some A/μs.

Also similar to SSCBs, but without integral isolation contacts, are solid-state overcurrent protectors, also called e-Fuses in the industry. They have breaking time in the same order as SSCB and, unlike conventional fuses, are resettable. Presently, such devices are considered only to provide supplementary overcurrent protection in accordance with NFPA 70.

Solid-state overcurrent protectors in the form of integrated circuits (IC), intended to be mounted on a printed wiring board, are covered by UL 2367 and by subclause G.9 of IEC 62368-1, being limited to a rating of 5 A (removal of such limitation in the new edition of UL 2367 is under discussion) and 250 VA. While there are no explicit restrictions on voltage, UL 2367 does not address requirements for devices rated above 60V.

Devices according to UL 2367 can be used for providing supplementary protection of electronic equipment, e.g. power supplies, in accordance with UL/IEC 62368-1 and UL/IEC UL 60950-1, and for ensuring selectivity. They can also provide protection for Class 2 circuits as defined in Article 725 of NFPA 70. Application for higher power distribution is currently not addressed.

Properties of different protection devices for overcurrent protection are shown in the table below.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

Table 7.2 Properties of different protection devices for overcurrent protection

	SSCB	**SCHCB** (hybrid breaker)	**EM CB TM**	**EM CB ETU**	**Fuse (power distribution)**	**Fuse (semiconductor protection)**	**e-fuse**
Reference standard	IEC 60947-10, UL 489I	IEC 60947-10, UL 489I	IEC 60947-2, UL 489	IEC 60947-2, UL 489	UL 248	UL 248-13	UL 2367
Suitable for branch protection	Yes	Yes	Yes	Yes	Yes (up to 600V)	Only suppl. protection	Only suppl. protection
Ground fault protection of equipment	Possible	Possible	Not possible	Possible	Not possible	Not possible	Not possible
Tripping Method	I threshold (dI/dt)	I threshold (dI/dt)	I threshold (M), I-t curve (T)	I threshold (dI/dt)	Energy	Energy	I threshold (de-sat)
low let-through current	++	+	-/- -	-	-	+/-	++
Interruption speed[a]	1-100s µs	0.1-1 ms	1-10s ms	1-10s ms	1-10s ms	0.1-1 ms (to be verified)	1-100s µs
Interruption capacity	unlimited	very high	very high	very high	very high	very high	unlimited
Power loss	--	+	++	++	+	+	--
Switching voltage	Medium	Medium	Low	Low	Low/Medium	High	Medium
Resettable	Yes	Yes	Yes	Yes	No	No	possible
Galvanic Isolation	Yes	Yes	Yes	Yes	With switch-disconnector	With switch-disconnector	No
Adjustable trip settings	Yes	Yes	Yes	Yes	No	No	No
Remote control (open)	Yes	Yes	Yes	Yes	No	No	(Yes)
Market availability	First products entering the market	First products entering the market	off-the -shelf	Very few products	off-the -shelf (up to 600V)	off-the -shelf	Only low currents

[a] for low-inductance short-circuits (high di/dt)

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

7.3. Protection against electric shock

Requirements regarding protection against electric shock in electrical installations are provided by IEC 60364-4-41.

Subclause 411.3.2 specifies requirements for automatic disconnection in case of fault, applicable also in DC. Requirements depend on the grounding system and the type of interlink converter.

In TN-S systems with isolated interlink converters, a protective device shall automatically disconnect and isolate all the supplies to the line conductor in the event of a low-impedance fault between the line conductor and an exposed conductive-part or a protective conductor in the circuit or equipment. According to Table 41.1 of IEC 60364-4-41, for voltage higher than 400 VDC, the disconnection time shall not exceed:

- 0.1 s, for final (branch) circuits with rated current up to 63 A, if including one or more socket-outlets, and up to 32 A if supplying only fixed connected equipment;
- 5 s, for distribution circuits (feeders) and for final circuits other than those mentioned above.

The same requirement applies to systems with non-isolated interlink converters, with grounded AC neutral and DC active parts ungrounded (see Section 6.2.1.3).

For IT systems, including both ungrounded and impedance-grounded systems, IEC 60364-4-41 does not mandate automatic disconnection on the occurrence of a first fault, provided a residual current monitor (RCM) or an insulation monitoring device (IMD) is used to signal the occurrence of such first fault with an audible and/or visual signal (IEC 60364-4-41, 411.6.3.1). Similarly, NFPA 70 does require capability to detect a first fault if the IT system is ungrounded (no intentional connection to ground through high impedance) in Section 250.167(A).

For grounding systems different from IT, subclause 411.3.3 of IEC 60364-4-41 also mandates use of a residual current protective device (RCD) with residual current of 30 mA to provide additional protection of AC sockets (receptacles) with rated current up to 32 A used by ordinary persons, for applications requiring special attention, e.g. in kitchen and bathrooms, lighting systems, outdoor installations and eV-charging. Additional protection for DC systems is under consideration.

Similarly, Section 210.8 of NFPA 70 identifies requirements for Ground-Fault Circuit Interrupter (GFCI) according to UL 943C (where applicable, other portions of the NEC like Section 422.5 may also identify further requirements), with tripping threshold not exceeding 6 mA, for the

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

protection of receptacles supplied in AC in certain locations or supplying specific appliances. GFCI requirements in NEC Section 210.8 also apply to receptacles required by NEC Section 210.63 for equipment requiring servicing, with rated voltage up to 150 VAC, in locations other than dwelling units. Additionally, NEC Section 215.9 permits the installation of GFCI in the feeder instead of the branch circuit. Further updates to the NEC and UL 943C will be needed to extend traditional AC GFCI (Class A according to NEC and UL 943C) requirements to address DC.

While today RCDs are rarely used in conventional AC data centers, their application in DC data centers could be considered, especially in case of TN-S supply to racks with 800 VDC internal vertical bus. However, such application requires special attention due to the risk of unwanted tripping that could impact system availability. Current/OS and UL are investigating technical specifications for DC RCDs.

Similarly, the use of IMDs for DC power systems in data centers could also be considered. IMDs continuously monitor the impedance between the live conductors and the ground. For ungrounded or high-impedance grounded systems, they are able to identify the first fault and warn the system controller/operator. Even for grounded systems, IMDs can help to detect the insulation degradation, such as the busbar and cable insulation degradation, and they can give warning before insulation breakdown and real damage to the hardware. However, the use of IMD requires careful technical evaluation: in systems with large parasitic and filter capacitance, the detection time of IMD on the first ground fault can be in the order of 15 minutes. In addition, IMDs alone do not provide selectivity to fault location, which will result in system wide reaction; locating faults require separate components and will further prolong detection time.

Additional considerations apply in case of maintenance, due to the presence of capacitors facing the DC bus: disconnecting the sources is not enough to ensure electrically safe working conditions, but also the capacitors must be discharged to a safe voltage level.

Safety requirements are stated in Article 360 of NFPA 70E. In case the voltage is 400 V or higher and stored energy exceeds the hazard threshold of 0.25 J, which is very likely in a DC power distribution system of a data center, a written procedure shall be developed stating all the required steps to discharge the capacitors and place the equipment into an electrically safe working condition. The safe voltage level is stated in NFPA 70E Section 360.3 and is 100 V for 100 J stored energy. Additional prescriptions and recommendations are provided in other sections of NFPA 70E Article 360 and Informative Annex R and in NEC Article 460, although not

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

directly applicable to the case of capacitors included as components in equipment. NEC Section 460.6 requires the residual voltage on a capacitor to be reduced to 50 V or less within 60 s after it is disconnected from the source.
Similarly, subclause 4.4.9 of IEC 61800-5-1 requires capacitors to be discharged to a voltage less than 60 V within 5 s after disconnection. If this requirement cannot be met, warning labels reporting the discharging time shall be applied.
Country-specific workplace safety regulations may set additional requirements.

Hazard resulting from large Y-capacitors shall also be considered (see Section 6.2.3.2).

7.4. Protection against thermal effects and fires

Insulation failures, ground faults and series-arc faults in power distribution systems, may evolve into large parallel arcs, resulting in electrical fires and arc-flash hazard. These risks need to be considered and further evaluated.
IEC 60364-4-42 sets out protective measures for locations where consequences of fire are severe, which include locations where people sleep (e.g. in dwellings or hotels), or intended for groups of persons, or museums, or others. Such measures are the same for AC and DC systems and include:

- for TN-S systems and systems with IT grounding and non-isolated interlink converter: installing an RCD (or a GFPE) with rated residual operating current not exceeding 300 mA (for non resistive loads);
- for systems with IT grounding and isolated interlink converter: installing an insulation monitoring device (IMD) monitoring the whole installation, and providing measures to eliminate a first earth fault with the shortest practicable delay.

While not mandatory, these measures should be considered for data centers due to their high economic value. Same considerations made in 7.3, about the use of RCD and IMD, apply.

For power distribution to industrial and data center loads, NFPA 70 indirectly addresses a specific type of ground faults, where those faults have low impedance to ground. This is done for AC through the requirements of several sections, including Section 210.13(A), and has recently been extended to DC through Section 210.13(B) and correlating updates to other sections as well, requiring branch-protection disconnecting means rated 1000A or more, installed in solidly grounded DC systems at voltage up to 1500 VDC, to provide ground- fault

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

protection of equipment (GFPE). For AC, this is supported by the UL 1053 listing and for DC this is under investigation by UL.
Additionally, Section 250.167 of NFPA 70 states that ground-fault detection is mandatory in ungrounded DC systems and permitted in grounded DC systems. Further investigations are needed about this requirement.

About series-arc faults, NFPA 70, Section 210.12(B) requires arc-fault protection by means of Arc-Fault Circuit Interrupters (AFCIs) in 120-volt, single-phase AC, 15- and 20-ampere branch circuits in dwelling units and similar locations, e.g. dormitories or guest-rooms in hotels. No similar requirement applies for commercial and industrial applications, neither in AC nor in DC. A special case is that of Photovoltaic (PV) DC systems, due to the harsh environmental conditions. In case of installation in or on buildings, Section 690.11 of NFPA 70 mandates arc-fault protection for voltage above 80 VDC; outline of Investigation UL 1699B specifies requirements for PV DC AFCIs up to 1500V DC. Equipment for detection and interruption of series-arc in PV systems are also addressed in IEC 63027:2023.
In DC data centers, arc-fault detection could be considered for circuits supplying racks with 800V vertical-bus. However, as the maturity of arc-fault detection in DC is relatively low, the risk of unwanted tripping is high. Considering the risk of arcs from hot swapping of modules, alternative mitigation measures are recommended, such as disabling swapping in case the power supply in the module is active, or automatically switching off power before swapping (semi-hot swapping).

7.5. Protection against voltage disturbances

Protection against voltage disturbances is addressed in Article 242 of NFPA 70 and IEC 60364-4-44.
In addition to voltage surges transmitted by the supply distribution grid, e.g. those of atmospheric origin, in DC systems special attention shall be paid to voltage surges generated by switching and breaking events, which can have significant magnitude.
All DC overcurrent protection devices, e.g. fuses, electromechanical circuit breakers, solid-state and hybrid circuit breakers, generate switching overvoltages which magnitude, duration and energy depend on the design of the device, the amplitude of the switched current, and the characteristics of the circuit. Protection devices based on semiconductors (SSCBs, SCHCBs and e-fuses) typically include surge protection devices (SPDs) to limit switching overvoltage

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

(see Figure above). IEC 60947-10 for SSCBs and SCHCBs, and IEC 60269-3 for semiconductor fuses require the manufacturer to declare in the product documentation the maximum switching overvoltage generated by the device.
Due to the low inductance and the relatively large capacitors on the DC bus, DC power systems in data centers are supposed to be relatively robust with regards to transient overvoltages. However, in designing the protection system, attention shall be paid to overvoltage coordination to ensure that overvoltage surges do not damage other electronic equipment in the system.

Another possible cause of the transient overvoltage on the DC bus is the insulation failure of the power converter between MVAC and the DC bus, such as the MV solid-state transformer (MV-SST). MV-SSTs are designed according to the insulation requirements in standards such as UL2877 (draft), IEC 62477-1 and IEC 62477-2. However, this is not enough to ensure that insulation failure will not happen during service life. Online monitoring of the insulation degradation for MV-SST is under development, but no mature solutions are available yet. Therefore, this topic needs special consideration.

7.6. Protection against arc-flash

There is the perception that the trend towards higher power densities in AI factories results in a significant increase of fire and arc-flash risks.
A single arc-flash incident can cause injury or death to personnel and destroy entire rows of expensive servers and other equipment, costing millions dollars.
Whenever possible, no work should be done on energized equipment. However, in data centers, typically operating on a 24/7/365 basis with strict uptime guarantees, turning off the power in a room or in a line of racks is impractical. As such, technicians may be required to perform maintenance and upgrade works on equipment while it is still energized.
For this reason, assessing and mitigating arc flash risks is crucial, in particular the adoption of adequate personal protection equipment (PPE), as described in NFPA 70E, protecting the workers thermal injuries resulting from the energy release or electric arcs (arc flash). This requires the Incident energy (IE_m) of the arc to be calculated.
New Section 110.16 of NFPA 70 2026 also requires permanent arc-flash hazard labels on distribution equipment likely to require servicing, or maintenance while energized. Additionally, Sections 240.67 and 240.87 of NFPA 70 require mitigation strategies be employed for circuits with overcurrent protective devices rated 1200 amperes or higher.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

As of today, there is no standard for calculation of IE_m in DC systems, equivalent to IEEE 1584 for AC. However, NFPA 70E Annex D.5 provides recognized methods based on different studies, one of which is the “Maximum Power Method” [D.5.1], known to give conservative IE_m values. The calculation formulas for this method are given below:

$$I_{arc} = 0.5 \times I_{bf}$$

$$IE_m = \frac{0.01 \times V_{sys} \times I_{arc} \times T_{arc}}{D^2}$$

Where:

- I_{arc} — arcing current (A)
- I_{bf} — system bolted fault current (A)
- IE_m — estimated DC arc-flash incident energy at maximum power point (cal/cm²)
- V_{sys} — system voltage (V)
- T_{arc} — arcing time (s)
- D — working distance (cm)

Where the arc is in a box or enclosure, the incident energy should be prudently multiplied by a safety factor, a reasonable value of which is 3.

A problem of this method is that the real behaviour of the fault current is not considered and a steady-state current I_{arc} is assumed for the total duration of the arc. Furthermore, it is necessary to clarify whether the value of the prospective short circuit-current I_{bf} to be considered is the peak value of the capacitor discharge, the steady-state fault current or a combination of those.

Even applying such conservative method, for a fault with short circuit-current I_{bf} = 100 kA, at 800 VDC, protected by means of a SSCB clearing the fault in 50 µs, the incident energy IE_m at 18 inch (45.7 cm) working distance is as low as 0.02 cal/cm² in “open air” (0.012 cal/cm² inside an enclosure).

In general, the graph in Figure 7.13, representing the IE_m calculated with the Maximum Power Method for different bolted fault current I_{bf}, in “open air” and “inside an enclosure” configuration, shows that the protection technologies in 7.2.5 and the selectivity strategies in 7.2.2 are sufficient to limit the arc duration so that the Incident energy remains below the threshold of 1.2 Cal/cm² that triggers the Arc Flash boundary for Personal Protection Equipment (PPE) according to NFPA 70E.

The results of simple calculations and simulations [Mar] confirm that relatively low levels of the available fault energy in the type of DC system used for data centers.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

More investigations are needed in the field of arc flash behavior in DC power systems, to confirm such considerations and develop industry-accepted methods for calculating incident energy and arc-flash boundaries.

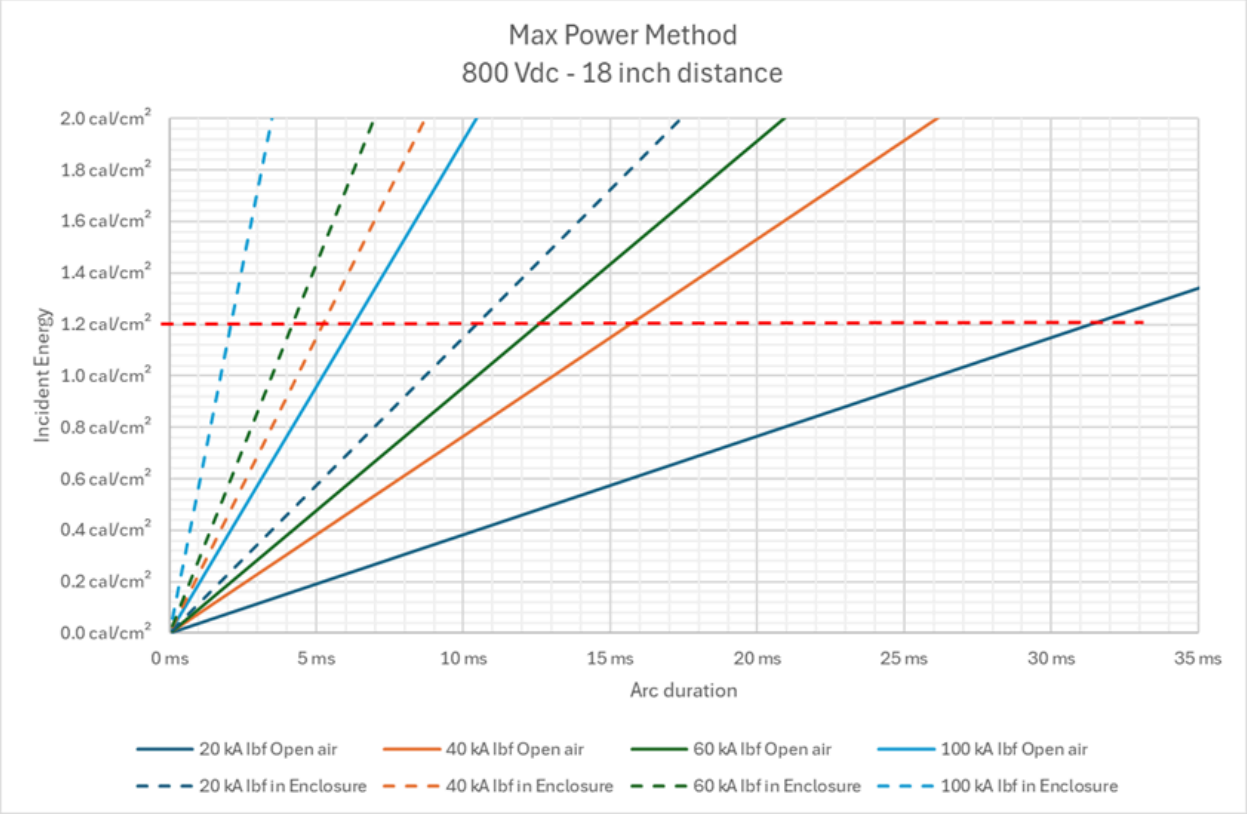


Figure 7.13. Incident Energy vs. Arc Duration

It is worth noticing that the Incident Energy is one hazard among others which results from an arc-flash event. In presence of capacitors, Article 360 of NFPA 70E defines additional safety measures related to the ‘arc blast” hazard: hearing protection is required where the energy stored in capacitors exceeds 100 J, which is quite likely to occur in data centers as energy storage is brought into the white space (with AC or DC power distribution). Whereas, the 122 kJ threshold, triggering the lung-protection boundary, is not likely to be exceeded.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

8. System Main Components

The “System Main Components” aims to list the principal power train components required to feed the AI Racks using DC power. For each component, relevant technology will be mentioned, and advantages and drawbacks will be discussed related to this specific application.
All the information shared in this part will remain high level unless it is deemed to be necessary to enter in more detail for clarification purposes.

Proposed Simplified Single Line Diagram:

The proposed SLD here is highly simplified, the only purpose is to facilitate the location of the different Main Components in the circuit.

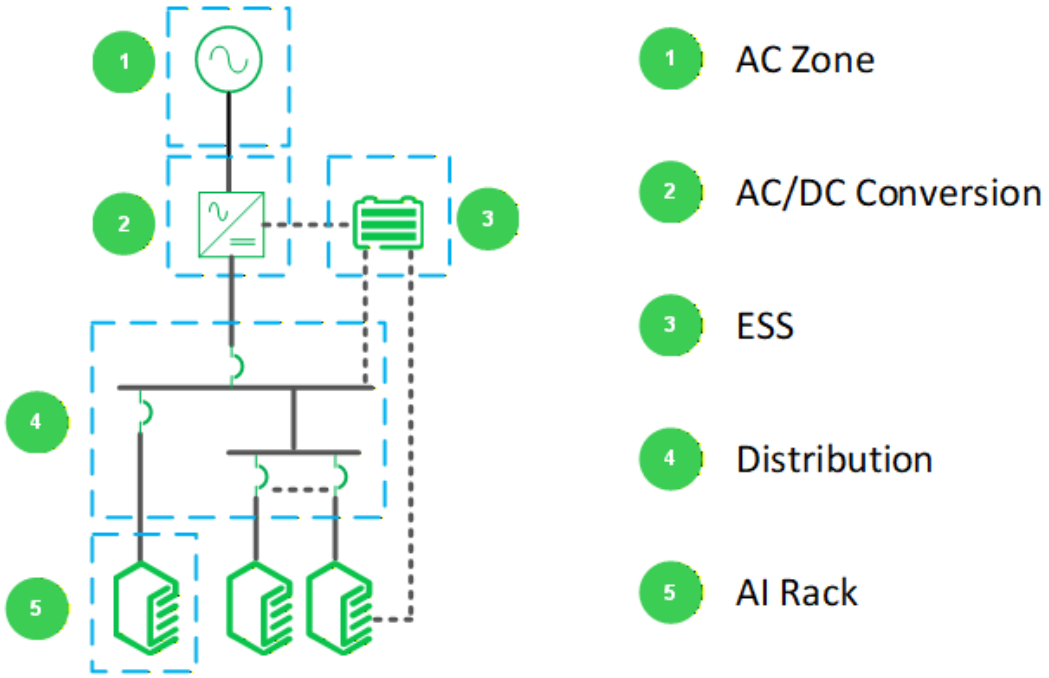


Figure 8.1. Simplified single line diagram for DC Power Distribution in Datacenter

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

The Main Components will be divided in several sections:

1. AC Zone

   The AC zone covers the requirements of the equipment upstream of the DC distribution side, such as connection to the grid, on-site generation, or even back-up generation.

2. AC/DC Conversion

   This section presents the different technologies available to convert AC into DC.

3. ESS

   This section highlights the different ESS technologies and where/how they can be used in the system.

4. Distribution

   This section covers the equipment required to distribute the power from the upstream conversion to the AI Rack.The protection equipment is not covered in this section as it has its own chapter.

5. AI Rack

   OCP has a dedicated Rack & Power work group, the information shared in this section are the ones which have been deemed impacting the DC power distribution requirements.

6. Miscellaneous

   In this section will be listed equipment which will play an important role in the overall architecture but have not been covered yet in the other section.

8.1. Utility Grid

For most data centers, the utility grid is their main source of power that is used to run the facility. Historically, data center load growth was relatively predictable and their impact on grid stability was minimal, grid operators were not concerned about data centers impacting the grid. However, as data center sizes continue to grow rapidly, driven by the accelerated development and deployment of AI workloads and reaching gigawatt scale, their impact on the grid reliability has become significant. Grid disturbances involving data centers have occurred across North America and Europe have prompted grid operators and transmission operators to institute requirements in their grid codes for data centers connecting to the grid. If these requirements are not met the data center will not be allowed to connect to the grid and energize. The sudden loss of load from a data center can cause high voltages and frequencies that can trigger grid protection systems worsening a disturbance. Therefore, data centers must be designed to meet

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

grid requirements to contribute to the overall reliability of the grid. These grid requirements are detailed in the subsections below.

Why are these requirements strict (in plain terms)? A modern AI data center is now large enough that the way it reacts to a brief grid fault affects the whole grid, not just the facility. Two things make it harder on the grid than ordinary load: it tends to draw more current as the voltage drops — deepening a dip instead of riding through it [17] — and the same load behaves worse on a smaller or weaker grid [18]. So the rules ask the facility to stay connected through short voltage and frequency disturbances and to bring its power back gradually, rather than disconnecting and then all reconnecting at once, which would hit the grid a second time [19].

8.1.1. Voltage Ride Through (VRT)

The data center must be able to withstand temporary voltage excursions for normally cleared faults by staying connected.

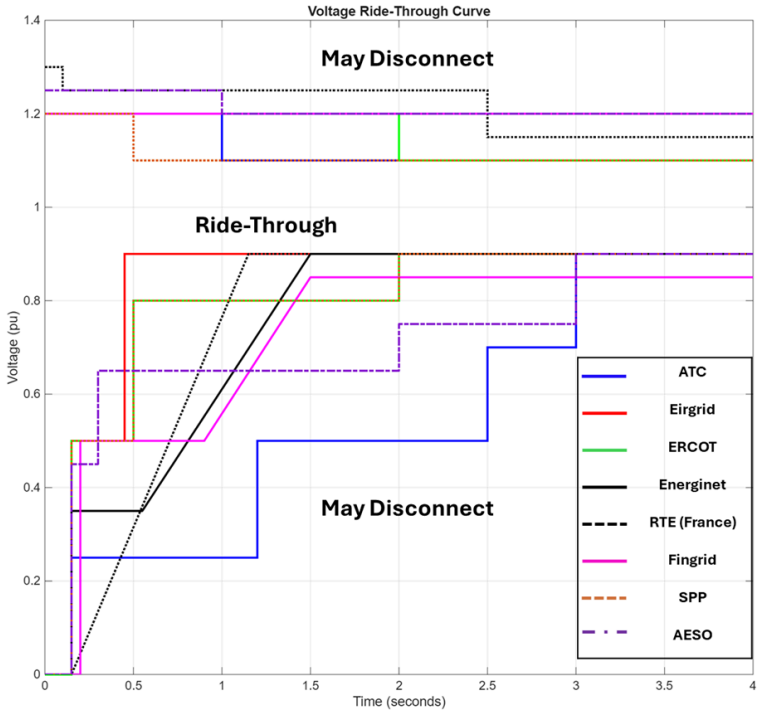


Figure 8.2. Voltage Ride-Through Curve

It is important to realize that when large voltage fluctuations occur, inrush currents and charging currents in the connected systems may delay the time that equipment further down the line will

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

see the full voltage, possibly up to 10s of seconds later. Keeping equipment energized near its nominal voltage will provide a better ride through but requires a local power system.

During a transmission fault the voltage at the point of interconnection (POI) falls until protection clears the fault. VRT requires the facility to remain connected and stable along a defined retained-voltage-versus-time envelope measured at the POI — the per-operator curves plotted in Figure 28. The reviewed requirements take two forms: North American operators specify a discrete minimum ride-through time per retained-voltage band (Table 8.1); European demand operators specify a continuous voltage-versus-time recovery trajectory under the ENTSO-E Demand Connection Code (Table 8.2).

Table 8.1. Ride-through time per retained-voltage band.

Retained voltage (p.u.)	ERCOT (US-Texas) [20]	SPP (US-Southwest) [21]	ATC (US) [22]	IESO‡ (CA-Ontario) [23]	AESO (CA-Alberta) [24]
V > 1.25	May ride through or trip	May ride through or trip	May ride through or trip	May trip	May ride through or trip
1.20 – 1.25				2.0 s	1.0 s
1.15 – 1.20	1.0 s	0.5 s	1.0 s		
1.10 – 1.15					1800 s
1.05 – 1.10	Continuous	Continuous	Continuous	1800 s	Continuous
0.90 – 1.05				Continuous	
0.80 – 0.90	2.0 s	2.0 s	3.0 s	2.0 s	3.0 s
0.75 – 0.80	0.5 s	0.5 s		0.5 s	
0.70 – 0.75					2.0 s
0.65 – 0.70			2.5 s		
0.50 – 0.65					0.3 s
0.45 – 0.50	0.25 s	0.15 s	1.2 s	0.25 s	
0.35 – 0.45					0.15 s
0.25 – 0.35	0.15 s				
0.20 – 0.25			0.15 s		
V < 0.20				0.16 s	

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

Table 8.2. Continuous voltage–time recovery trajectories.

Authority / region	Status	Deepest-sag ride-through	Recovery floor (voltage must stay above)	Notes
EirGrid / SONI (Ireland) [25]	Grid Code mod. recommended Apr 2026; pending CRU	0 p.u. for ≥ 0.15 s	≥ 0.50 p.u. at 0.15 s → ≥ 0.90 p.u. by 0.45 s	—
Energinet (Denmark) — ≥ 200 MW [26]	In effect 2024	0 p.u. for ≥ 0.15 s	≥ 0.35 p.u. to 0.55 s → ≥ 0.90 p.u. by 1.50 s	Current blocking allowed in Area B
RTE (France) [27]	In effect 2024	0.05 p.u. for ≥ 0.15 s	linear recovery to ≥ 0.90 p.u. by 1.15 s	PE-interfaced consumption
Fingrid (Finland) [28]	In effect 2026	0 p.u. for ≥ 0.20 s	≥ 0.50 p.u. (0.2–1.0 s) → ≥ 0.85 by 1.5 s → ≥ 0.90 by 10 s	8 × 100 ms faults / 90 s; < 0.50 may cease

Two forms of requirement. The North American operators express discrete retained-voltage bands each with a minimum ride-through time and handle active-power restoration separately as PFAPR (§8.1.3). The European operators specify a continuous voltage-versus-time recovery trajectory which cannot be converted to discrete bands without interpolating values not stated in the source. The split is one of the regulatory forms; it holds across the operators reviewed here but should not be assumed universal, and the band form is itself not monolithic (IESO uses site-specific Vmin/Vmax).

8.1.2. Frequency Ride-Through (FRT)

The data center must be able to withstand temporary frequency excursions for normally cleared faults by staying connected.

Following a generation-load imbalance, system frequency deviates from nominal. FRT requires the facility to remain connected within defined frequency-versus-time bands so its response does not intensify the event. The 60 Hz operators are compared in Table 8.3; EirGrid/SONI (50 Hz) is given in Table 8.4.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

Table 8.3. 60 Hz frequency ride-through envelopes. ERCOT/SPP/IESO per the existing §5.8.3 contribution; AESO per Guide Table 6-3 (verified).

Frequency (Hz)	ERCOT [20]	SPP [21]	IESO [23]	AESO [24]
f > 63.0	May ride through or trip	May trip	May trip	May ride through or trip
61.8 < f ≤ 63.0	≥ 299 s			
61.2 < f ≤ 61.8		≥ 299 s	≥ 299 s	≥ 299 s
58.8 ≤ f ≤ 61.2	Continuous	Continuous	Continuous	Continuous
57.5 ≤ f < 58.8	≥ 299 s	≥ 299 s	≥ 10 s	≥ 299 s
57.0 ≤ f < 57.5	May ride through or trip	May trip		
f < 57.0			May trip	May ride through or trip

Table 8.4. EirGrid/SONI proposed a 50 Hz frequency operating envelope. RoCoF (verified): EirGrid/SONI ± 1 Hz/s over a rolling 500 ms; AESO ≤ 5 Hz/s over a rolling 500 ms. AESO additionally specifies a phase-angle limit ≤ 25° and V/Hz ride-through (≥ 1.1 p.u. for 45 s; 1.18 p.u. for 2 s)

Frequency (Hz)	EirGrid/SONI — minimum ride-through time without disconnection
f < 47.0	No minimum time
47.0 ≤ f < 47.5	≥ 20 s
47.5 ≤ f < 49.0	≥ 90 min
49.0 ≤ f ≤ 51.0	Unlimited
51.0 < f ≤ 52.0	≥ 90 min
f > 52.0	No minimum time

8.1.3. Post Fault Active Power Recovery (PFAPR)

The data center must be able to expeditiously reconnect to the grid after the disturbance is cleared. For example some grid operators are requiring that data centers reconnect and return to roughly 90-100% of pre-fault power consumption within one second. Therefore data centers must be able to program settings on how quickly the data center reconnects to the grid and returns to a specified pre-fault power consumption level in the event of a fault. Also in this case the delays occurring due to charging currents and inrush currents that delay the reestablishment

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

of active power to pre fault levels will require strategies like near the POI power buffers and precharging/islanding.

PFAPR is assessed at the facility POI: after voltage recovers, the facility must restore active power to a target fraction of its pre-disturbance level within a defined time. North American operators are compared in Table 8.5 and European operators in Table 8.6. Rack-level PSU fast-resume and BBU handoff are enabling functions, not the compliance measurement point.

Table 8.5. Post-fault active-power recovery — North American operators.

Operator	Recovery target	Recovery time (after POI V ≥ 0.90 p.u.)	Qualification
ERCOT [20]	≥ 90%	≤ 2 s	HVAC/cooling load may trip below 0.35 p.u.; computational load must recover
SPP [21]	≥ 90%	≤ 1 s	Deep sag (< 0.50 p.u.) may transfer load to internal sources
IESO [23]	100%	≤ 1 s (default; configurable)	No normal-operation ramp limit may delay post-fault recovery
AESO [24]	100%	≤ 1 s (default; slower in weak grids)	No momentary cessation within the ride-through envelope

Table 8.6. Post-fault active-power recovery — European operators

Operator	Recovery target	Recovery time / rate	Qualification
EirGrid / SONI [25]	≥ 90%	within 500 ms of voltage recovering to 0.9 p.u.	Measured at the connection point
Energinet [26]	staged ≥ 80% / ≥ 90% / ≥ 95%	within 5 s / 20 s / 30 s after return to the continuous range	Category-7 demand (≥ 200 MW)
RTE [27]	≥ 90%	ramp 0.5 Pmax/s after V > 0.9 Un held 300 ms	PE-interfaced consumption
Fingrid [28]	≥ 90%	within ~1 s after V → 0.9 p.u.	—

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

8.1.4. Protection System Coordination

The data center protection systems logic must consider timing needs for high speed reclosing and transmission breaker operations and coordinate with the grid protection to minimize misoperations. Any other control logic that is impacted by grid disturbance should also be coordinated to minimize misoperations.

8.1.5. Ramp Management

On the grid, demand and generation must remain balanced at all times. If demand or generation rises or falls too quickly then the grid may exhaust its power reserves and be unable to manage frequency threatening grid reliability. To avoid contributing to this issue, data centers must manage their active power consumption ramp rates by controlling how quickly the data center load ramps up and down. Ramp control may be executed by means of local generation or energy storage systems, creating a buffer between the DC power requirements and the load taken from the utility.

8.1.6. Oscillation Management

Repeated power oscillations from a data center can be a source for forced oscillations on the grid which can damage generators. Data centers focusing on AI training have a higher oscillation potential including sub-synchronous oscillations than previous data center designs. Data centers must manage and avoid repeated power fluctuations to manage the potential for oscillations.

8.1.7. Power Quality

Data centers must meet power quality requirements to avoid impacting other customers and causing reliability problems.

8.1.8. Monitoring Methods

Data centers must be able to monitor the performance of the data center as a facility.

8.1.9. Reactive Power and Voltage control

Connecting large data centers to the grid can have a significant impact on voltage and reactive power flows. It is expected that the data center actively participates in voltage management by

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

providing reactive power within a predefined range (typically ±0.95 power factor) and can be coordinated with the setpoint from the connecting Transmission System Operator.

8.1.10. Validated Grid Models

Grid planning engineers rely on accurate electrical models that simulate data center behavior. Data center owners must develop and provide high-fidelity dynamic models in the positive-sequence (e.g., PSS/E) and EMT (e.g., PSCAD) domains that accurately reflect the facility dynamic responses.

8.2. On-Site Generation

As AI-driven applications scale, technology companies have traditionally relied on the local grid to power their data centers. But with soaring demand, the grid is struggling to keep up, facing capacity and transmission constraints. As a result, data center developers—who typically focus on facility design and optimization— are increasingly looking for alternative power solutions that are reliable and scalable.

On-site generation can function as behind-the-meter (BTM), island microgrid solutions, enabling data centers to operate independently where grid infrastructure is insufficient or unavailable.

In locations where grid operators require onsite power to contribute to the electricity market, these technologies can also run in parallel with the grid as dispatchable resources, providing additional flexibility.

In the case of on-site power generation, the applicable requirements for grid interconnection when grid is the source of power, and described in the previous chapter, are no longer needed.

However additional challenges apply, especially resulting from the following:

- A direct influence of AC/DC conversion harmonics onto the microgrid power quality: higher reactive power, losses, frequency and voltage stability;
- The reduced ability of the on-site generation to immediately react to the AI workload oscillations, given the more limited generation capacity compared to having the whole grid as a source.

8.3. AC/DC Conversion

With higher power densities and a DC distribution system, facility level AC/DC conversion is needed to supply the DC distribution system. By using a centralized conversion system, flexible DC and bidirectional power flow capabilities are enabled allowing for point-to-point power

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

routing, improved fault handling, and seamless integration with battery energy storage systems. Several systems can be used for AC/DC conversion including uncontrolled rectifier “x” pulses, LV Interlink Converters, MV Interlink Converters, and Solid State Transformers.

The purpose of this section is to present a technical evaluation of different features and technologies available to convert AC into DC with high level pros and cons. Future versions of this section will include more analysis around maintenance, cost, integration or efficiency.

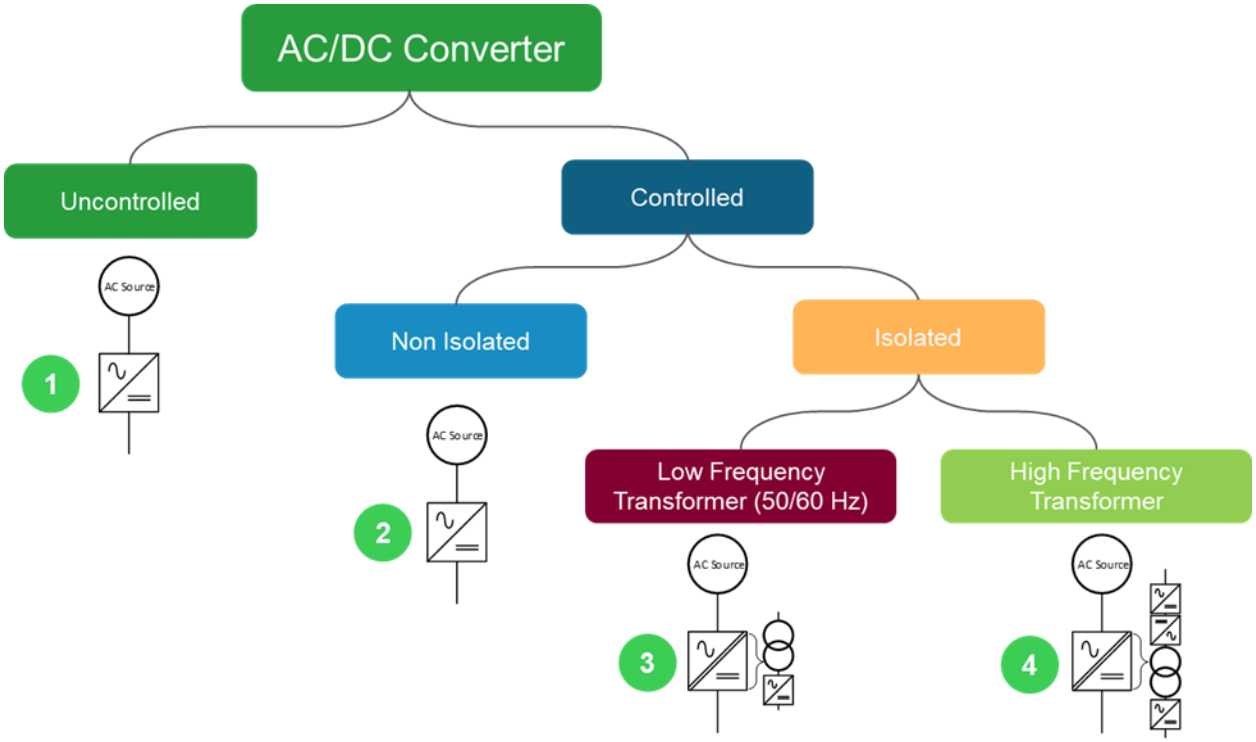


Figure 8.3. AC/DC Converter Block Diagram

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

Uncontrolled

- Vdc output is dependent on Vac input, any fluctuation on Vac will impact Vdc
- Vdc is always >Vac
- Unidirectional only

Controlled

- Vdc output is independent of Vac input (to the limit of the converter topology)
- Vdc is always >Vac for “boost type”
- Vdc is always <Vac for “buck type”
- Vdc can be > or <Vac for “boost/buck type”
- Unidirectional or bidirectional

Non Isolated

- No galvanic isolation between ac and dc side
- Grounding configurations between ac and dc are dependent

Isolated

- Galvanic isolation between ac and dc side
- Grounding configurations between ac and dc are independent

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

Low Frequency Transformer (50/60 Hz)

- Galvanic isolation is provided by a low frequency transformer (50/60 Hz)

High Frequency Transformer

- Galvanic isolation is provided by a high frequency transformer

1

“x” pulses Rectifier type (6,12,24,36…) using semiconductors such as diodes.

- Pros
  - Mature and reliable technology
  - Relatively less sensitive to fault current
- Cons
  - Energy control not possible (e.g droop)
  - May required harmonics mitigation devices
  - Unidirectional Only

While potentially a valuable solution for DC Datacenter (similar technology has been used in the previous project in the world), the lack of output voltage regulation will affect the stability and power quality of the system for AI type of load profile, which makes this solution hard to implement (it would require high complexity of controls).

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

2

“boost” type using semiconductor such as IGBT, MOSFET. Sometimes called “Active Front End” ‘AFE) or Interlink Converter.

- Pros
  - LVAC input only to provide LVDC output
  - Mature and reliable technology
  - Very efficient (reduced semiconductor switched device)
- Cons
  - Grounding configuration must be engineered to avoid incompatibility between ac and dc side
  - Sensitive to fault current, fast tripping time required to protect power electronic components

3

Similar to number 2 adding a low frequency transformer at its input

- Pros (in addition of number 2)
  - Provide galvanic isolation
  - Can take LVAC or MVAC as input
- Cons (in addition of number 2)
  - Solution can be bulky for very high power due to the LF transformer

These types of converter (2 and 3) are widely used in the industry today (Solar, BESS, Drive e.g). This solution is a good fit to provide stable and quality power to the AI rack until the integration (footprint, weight, transportation) becomes a constraint which can lead to the use of converter type number 4.

4

Several topologies exist, using LVAC as input (e.g LV rectifier+dual active bridge) or using MVAC as input also known as Solid State Transformer (SST*)

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

- Pros
    - Provide galvanic isolation
    - Can take LVAC or MVAC as input
    - Reduce size (footprint and weight gain)
    - Some topologies allow control of the current during fault condition hence the use of slower protection devices.
- Cons (in addition of number 2)
    - “young” technology compared to the other ones
    - Reliability at scale to be proven (high number of switching devices)
    - Efficiency at scale to be proven

*A Solid-State Transformer (SST) is defined as a megawatt-scale power electronic converter designed to replace conventional 50/60Hz line-frequency transformer and rectifier units (isolated or non-isolated). The distinguishing features of SST include medium-to-high-frequency magnetic components/transformers (MFT/HFT), enabled by Wide-Bandgap (WBG) semiconductors such as Silicon Carbide (SiC) or Gallium Nitride (GaN).

Adding an ESS+Control system on the output of the AC/DC converter and packaging it together makes a single equipment which can be called DC UPS (DC output instead of traditional AC).

Converters can also be further classified by input and output configuration:

1. Single-Input Single-Output (SISO) SST
    - One AC input and one regulated DC output
2. Multi-Input Multi-Output (MIMO) / Multiport SST
    - Multiple AC and/or DC inputs and multiple regulated DC outputs

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

There is no standardized “fault current behavior” today, in addition to being physically linked to the converter topology, controls (manufacturer based) also play a role.

Fault current behavior can be split into 2 categories:

1. Fault feeding (converter type 1,2 and 3)

In this category, the fault current from the AC source to the DC distribution cannot be limited. Converter type 1 can usually withstand fault current for a short duration, however, converter type 2 and 3, because of the semiconductor switches, cannot withstand fault current coming from the AC source. These types of converter usually have self fast protection (e.g fast fuse or semiconductor switch) embedded which will limit the fault current in duration and or in peak value.

2. Fault Limiting (converter type 4)

These types of converter, thanks to the inherent topology, can control and limit the fault current from the AC source to the DC distribution. It can limit the value to a small multiple of the nominal current (e.g 1.5x, 2x) and limit in duration (e.g 100ms, 200ms).
It is worth noticing that in any case, there could be a short surge current when the fault happens because of the exposed capacitor of the converter on the DC side.
Principal Standards which covers AC/DC Conversion technologies:

- UL 1741 → Standard covering Converters
- UL 2877 → Outline of Investigation which covers SST
- IEC 62477-1 and 2 → Standard covering Converters
- IEEE 1547 → Standard for grid interconnection

Refer to the Codes & Standards chapter for more details.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

8.4. ESS

Energy storage systems (ESS) are becoming critical components of modern data center infrastructure. They go beyond the traditional role of an Uninterruptible Power Supply (UPS) by offering fast stabilization for AI-driven load swings, reliable ride-through during grid disturbances, cost optimization, renewable integration, and grid services.

The figure here under represents applications by the time frame the storage is expected to be active: milliseconds for local power quality, seconds for rapid stabilization and ride-through, minutes for UPS bridging and peak shaving, and hours for energy arbitrage, renewable shifting, and extended backup:

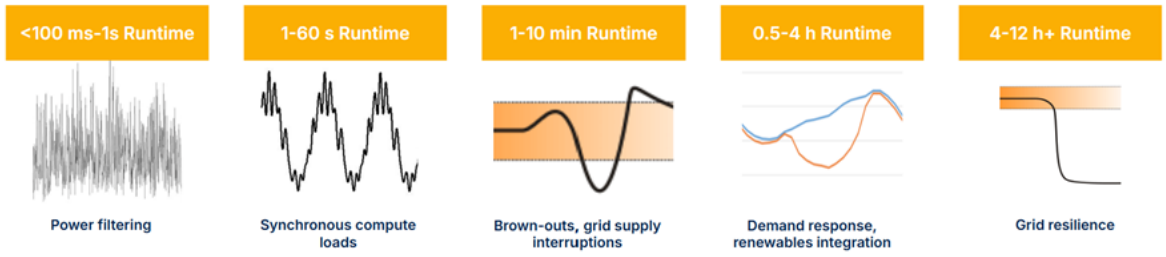


Figure 8.4. Power Equipment Applications

The key functions of ESS in a data center:

1. Smoothing power swings (milliseconds to seconds)

AI workloads present high power fluctuation, ramping power from 30% to 100% in milliseconds [Choukse et al. 2025]. These ramps and spikes are characterized by different dynamic elements occurring at different timescales (milliseconds, seconds, minutes), thus stressing equipment and degrading overall power quality. Without mitigation, the result can be voltage sags, oscillations, threaten broader grid stability, and risk tripping or damaging nearby generation and other upstream equipment. Energy storage systems address these dynamics by rapidly absorbing or injecting power across the relevant timescales, smoothing ramps and spikes, damping oscillations, and maintaining stable voltage and frequency through coordinated controls.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

2. Providing backup power during grid disturbance (seconds to minutes)

This is the most fundamental function. Data centers require uninterrupted power for their mission-critical servers and other systems.

Traditionally, backup of the facility is by UPS, and battery backup units (BBUs) are also deployed in or close to IT racks (sometimes in conjunction with supercaps or alternatively implemented with integrating both BBU and cap-like performance in a battery capacitor). They ensure that the data center remains operational during grid failures or until backup generators can be brought online. The timescales required for these types of operations are generally considered to be minutes. Alternative solutions could be used, for example, ESS with extended runtime to even replace generators entirely in some cases. By using batteries (for example) instead of diesel (or gas) generators for backup, data centers can achieve a zero-emission backup solution, supporting sustainability goals and providing a “friendlier” solution than noisy generators (especially in densified environments). The size and cost of the battery will be directly proportional to the amount of time that a given power level is desired to be provided.

Moreover, during Post-fault active power recovery the data center must be able to quickly reconnect to the grid after a disturbance is cleared. Some grid operators require large loads to return to roughly 90–100% of pre-fault power consumption within about one second. ESS enables this by “mimicking” the load for the grid, absorbing power during the reconnection ramp to keep the facility’s net import near its pre-fault level until the grid stabilizes. This requires the charge level on the batteries to be carefully controlled, as the battery (or battery capacitor) will need to be able to absorb a certain amount of power to comply with the PFARP and at the same time allow for covering power loss during a short event.

3. Grid Services and Energy Optimization (minutes to hours)

In many regions there is a paid market for grid services (e.g., frequency regulation and reserves), while in others the emphasis is on strict grid-code compliance; ESS can support both, optimizing economics where markets exist and ensuring reliable, compliant operation everywhere.

- Energy arbitrage: store low-cost electricity and discharge when prices are high to optimize energy costs.
- Distribution efficiency: smooth power profiles to keep electrical and computing systems in efficient operating regions and reduce losses. **NOTE**: performing this function at/near the rack level can enable very significant reduction of upstream

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

BESS that must be greatly oversized in capacity to account for the cycle life of highly-dynamic load smoothing.

- Peak shaving and capacity reduction: rapidly discharge during demand spikes to power demand charges and, by limiting instantaneous draw, reduce maximum contracted capacity and ease interconnection.
- Grid stabilization and ramp compliance: inject or absorb power to stabilize frequency, participate in load shaving, and meet ramp-rate requirements set by grid operators.
- Integration with renewable energy: pair with on-site solar or wind to store excess clean generation and use it later, maximizing renewable utilization and supporting decarbonization targets.
- Reactive power and voltage control: provide reactive power and voltage support when required by the grid.

8.4.1 Technology

The landscape of energy storage for data centers is rapidly evolving beyond the traditional UPS, with different technologies offering distinct trade-offs in terms of duration, energy/power density, safety, typical connection point and cost.

Proximity is relevant to how those needs and characteristics are provided for. Figure 8.4 below illustrates commonly considered connection points in the data center architecture where they may be integrated. ESS that is located closer to the IT rack will be tasked with the highest response rate, high cycle count, shorter duration dynamics. While ESS that is located further upstream or at the utility interconnect will deal with applications required for grid compliance or energy management. The utility-scale BESS generally applied for the slower dynamics and longer duration, whereas E-STATCOM for the fast dynamics and short duration, such as load swings and low-voltage fault ride through. The configurations and voltage ratings that are feasible for a given ESS technology must be suitable for the voltage level at the point of connection for that ESS within the power distribution system - and/or make use of a DC-DC converter. The core challenge is how to solve the different load dynamic challenges at different levels of the power distribution system with a coordinated and optimized set of ESS technologies. Table 8.5 is provided to give a summary of the nominal capabilities and strengths of different energy storage technologies to provide context for how they may be utilized or combined to provide for the overall load dynamic requirements.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

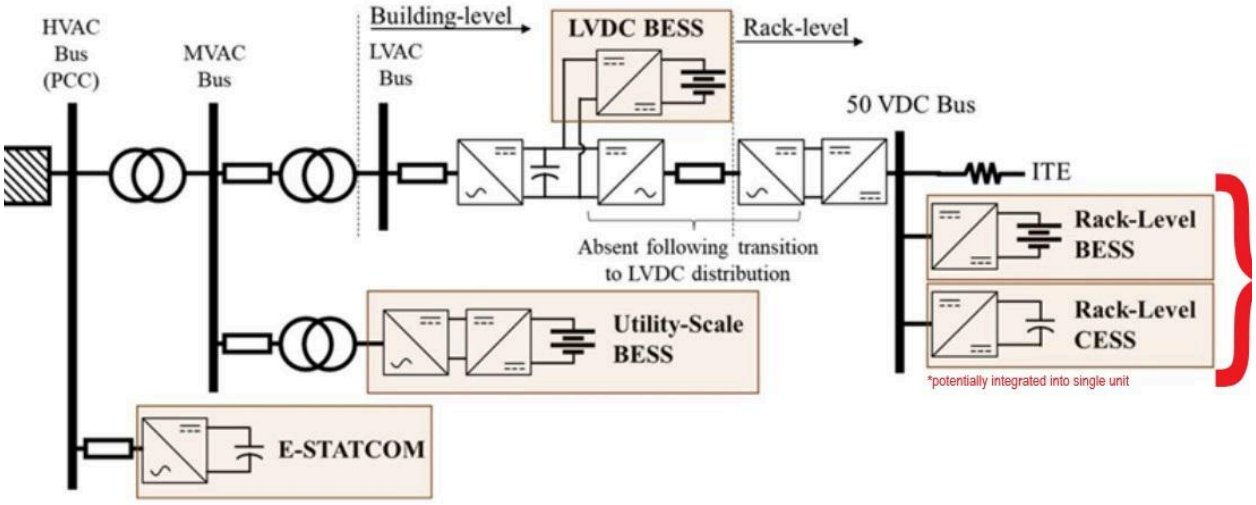


Figure 8.4. Electromagnetic Transient Modeling of Large Data Centers for Grid-Level Studies

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

Here is a breakdown of the primary energy storage technologies to be considered in data centers:

1. Electrochemical energy storage

   The conversion and storage of energy are through electrochemical reactions. It includes lead-acid battery, lithium-ion battery, metal air battery, flow battery, sodium-ion battery, battery capacitors and so on. The ESS based on batteries typically includes battery modules (may inc. power electronics), thermal management system, and battery management system.

2. Mechanical energy storage

   Mechanical energy storage systems store energy in the form of potential or kinetic energy. Unlike electrochemical energy storage, these systems use physical forces and motion to capture and release energy. It includes flywheels, compressed air energy storage, and so on.

3. Electrical and magnetical energy storage

   Electrical magnetic energy storage primarily refers to Superconducting Magnetic Energy Storage (SMES), they store energy directly in a magnetic field created by the flow of direct current (DC) in a superconducting coil. Conventional capacitors fall into electrostatic energy storage, they store energy directly in an electric field created by the separation of static electric charges. Supercapacitors function via a hybrid mechanism, by using an electrolyte and store charge in an Electric Double Layer (EDLC) formed at the interface between the electrolyte and the electrode (usually porous carbon).

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

Table 8.5: Main energy storage technologies considered for data center

Feature	Ceramic Caps (MLCC)	Electrolytic Caps (Al, incl. polymer/ hybrid)	Film Caps (PP/PET DC-link)	Supercaps (EDLC / hybrid)	Battery Cap	Li-ion – LTO XNO NTO	Na-Ion – NVPF	Na-Ion – Prussian White	Li-ion – NMC 811	Li-ion – LFP	Lead-Acid Batteries (VRLA/ Flooded)	Al-Air Batteries
Primary Role	Point-of-load HF decoup, edge suppress	Bulk decoup, hold-up energy	DC-link ripple handling, damping	Bridge power, pulse buffering, power quality	High-power buffer/ ride-through, fast charge/ discharge buffering, pulse peak shaving, SD-ESS	UPS, SD-ESS, fast charge/ discharge buffering	High-power buffer / ride through	BESS	BESS, automotive	BESS, automotive	UPS ride-through + backup energy	BESS
Storage Mechanism	Dielectric (ceramic)	Dielectric (electrostatic)	Dielectric (polymer film)	Electrochem Cap	Capacitive & Faradaic	Electrochem	Electrochem	Electrochem	Electrochem	Electrochem	Electrochem	Electrochem
Typical Runtime	ns → ms	ns → ~100 ms	µs → 10s ms	Sub-s → s	Sub-s → mins	s → mins	s → mins	mins → hrs	mins → hrs	mins → hrs	mins → hrs	hrs → days
Energy Density	Very Low	Low	Very Low	Very Low	Medium	Low–medium	Low	Medium	Very High	Medium	Low	High
Power Density	Extremely High	High	Very High	Very High	Very High	Very High	Very High	Medium-high	High	Medium	Medium	Medium
Cycle Life	Extremely High	High (if lightly stressed)	Very High	Extremely High (100k–1M+)	Extremely High (100k–1M+)	Very High (10k–20k+)	High-Very High	Medium-high	Medium–high (thousands)	High (1ks to 10k+)	Low (100s)	N/A
Lifespan	Long (but can crack)	Medium	Long (often 15–25+ yrs)	Long (10–20+ yrs)	Long	Long (10–20 yrs)	Long	Long	Medium–long (8–15 yrs)	Long (10–20 yrs)	Short (3–6 yrs)	Long (10-20 yrs)
Thermal / Safety Risk	Very Low	Low	Very Low	Low	Low	Low	Medium	Medium	Medium–high	Low–medium	Low	Low
Maintenance	Very Low	Low (but planned replacement)	Very Low	Low (monitor balancing/ health)	Low-medium	Low-medium	Low-medium	Low-medium	Low-medium	Low-medium	Medium	Low
Response Time Band	sub-ns → ms	ns → 100 ms	10s of ns → ~10s of ms	µs → s	µs → s	sub-ms → s	ms → s	ms → s	ms → s	ms → s	ms → s	ms → s
Voltage Rating Practicality (ie for 800 V bus)	Low	Medium-low	Medium-high	Low-medium	High	High	High	High	High	High	High	High

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

8.4.2 ESS Safety

ESS safety is a complex subject. The intent here is to provide a high level guide to ensure awareness of the most critical aspects and some understanding of why they should be considered. Most any ESS technology or medium poses at least some kind of thermal risk. Any ESS that utilizes an electrolyte consisting of an organic carbonate solvent will, in a failure scenario, produce an offgas that poses a risk for flammability and/or explosion. This summary will discuss primarily in the context of lithium-ion batteries, most electrochemical storage technologies (particularly those using organic carbonate solvent electrolytes) will be more similar than different. However, it is important to note that some lithium-ion batteries and battery capacitors decrease or completely mitigate some of these risks through the use of different materials (aqueous electrolytes, ceramic separators, oxide-based anodes), making them safer and more suitable for high-risk or mission critical applications. Capacitors typically present lower or no fire risk (ie ceramic), but technologies using electrolytes (electrolytic, supercapacitors) do constitute thermal risk. In most cases, thermal events are characterized by very low likelihood of occurrence, but the impacts of an event can be so severe that the overall risk profile demands attention. The risks and results of thermal events can be managed and mitigated with existing engineering tools and technologies. But variations exist between different technologies and form factors such that the mitigating solutions must be fit for purpose. This poses a challenge for ongoing efforts towards standardization and for early stage project scoping. In addition, human error is often found to be the failure mode behind many historical ESS safety events.

Standards and best practices for large outdoor ESS installations do exist with years of history. However, the prospect of a significant amount of energy storage capacity distributed throughout an indoor space poses new and different challenges. The following summaries discuss general aspects of ESS safety for consideration, but primarily in the context of challenges and needs for those indoor installations.

1. Thermal runaway: A self-accelerating thermal event that is characterized by exothermic processes. Any medium or technology that stores chemical energy is capable of and at risk of some degree of a thermal event. However, the amount of thermal energy released and the extent of abuse that is required before forcing a cell into thermal runaway certainly varies in different chemistries. However, suggestions that some battery technologies are completely incapable of a thermal event should be considered with skepticism. Some technologies may produce benign responses to certain failure modes,

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

but still be at risk of significant thermal events in response to other failure modes. Objective testing and valid data is critical for assessment of the true risks, capabilities and requirements. Generally, a single thermal event (i.e. in a single cell) is of relatively mild consequence to the overall facility. The primary threat and challenge is the likely scenario that this single cell event will propagate to other cells, and thus spread exponentially. This is a critical aspect of ESS system design and the requirements are particular to the chemistry and module/system design and form factor. Evaluating thermal and propagation behaviors is a core component of UL 9540a testing.

2. Electrical: a great deal of consideration is given to thermal aspects of ESS safety. However, practically speaking, the vast majority of risk and real incidences are related to electrical shock and arc flash. Batteries present voltages often over 1,000 VDC and, critically, are always live. Proper precaution, processes and equipment for aspects such as disconnect and isolation, particularly during installation and service, are critical to ensuring safety.
3. Ventilation: In the case of a thermal event, offgas will be released that is both flammable and explosive. Managing and ducting this gas to avoid an explosion is a key challenge - particularly for indoor installations. There are many different approaches that can be considered for mitigating risks - appropriateness depends on the technology, form factor, and the space within which it is installed. Ventilation poses one of the most substantial challenges with regard to infrastructure and available space, and is one of the required safety systems that is often forgotten until the end.
4. Fire suppression: Generally speaking, putting out a lithium ion battery fire is difficult and access directly to the cell in thermal runaway is typically not possible. Thus, ‘extinguishing the fire’ is not a typical objective. A key objective is to limit propagation, as described earlier; and for this, removal of heat is critical. Many different fire suppression mediums exist, and each accels in certain areas and may underperform in others. Gaseous mediums excel at extinguishing flames, but not at removing heat and must take into account interactions with ventilation. Water is ideal for removing heat but produces short circuit risks. Many others face challenges of physical installation feasibility. The system should be selected and engineered appropriately for the ESS technology and form factor that is being installed.
5. Codes and Standards: ESS codes and standards are continually evolving. The differences between technologies and formulations makes it particularly difficult for ‘one

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

size fits all' safety requirements. For commercial installations in the US, the most relevant standards which will govern ESS installation are UL 1973, UL 9540, UL 9540a and NFPA 855. UL 1973 has been the longest standing ESS fire safety standard. UL 9540 addresses requirements for the ESS product. UL 9540a presents a testing regime which elicits the behaviors of what happens when a given ESS fails. NFPA 855 governs site requirements for the installation of an ESS at a location, such as sizing, proximity, spacing. Given the evolving nature of the ESS standards landscape, in all cases the Authority Having Jurisdiction (AHJ) will be a driving force for what is required.

8.5. Distribution

Depending on the architecture, layout, and power density, there might be different means to safely and efficiently transfer power between the point of AC/DC conversion to the AI Rack.
This section introduces some of these elements:

- LVDC equipment such as Switchboard or Panelboard
- Conductors such as cables, busway
- Connectors which connect each element to each other

Each of these equipment must be certified/listed for DC use.

As in the previous chapter, this section includes high level information on the different available technologies, and highlights the challenge and benefits which each raises in a high power DC distribution system.
As a reminder, circuit protection which would fall under this section, is covered in Section 7.

- LVDC Equipment

LVDC equipment includes Switchboard and Panelboard (including RPP, Remote Power Panel). There are no specific requirements beside the fact that it must be designed according to the selected architecture, 2 or 3 wires, the rated voltage and current, and the protection device installed must open all ungrounded polarities (refer to chapter on polarity and earthing as well as protection device).
The principal applicable standards are:

- UL 891 → for Switchboard
- UL 67 → for Panelboard
- IEC 61439-2 → for Switchboard and Panelboard

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

Note: as of when this whitepaper will be published, UL 891 and 67 does not allow rated voltage > 1000 Vdc.

- Conductors

Conductors can be divided into different groups and technology.

“Traditional” Cable or Busway are a natural choice in the industry to transfer power to the load, in this case, the AI Rack.

While a valuable option for the DC distribution part for low to medium power (tens to hundred Watts, or low megawatts), this solution may introduce challenges when it comes to very high power (few megawatts). Especially around the infrastructure which would need to supper the space and weight of multi conductors or large busbar.

For very high power distribution, another solution possible is the use of superconductor technology. Superconductors are unique materials that exhibit zero electrical resistance when cooled below a defined operating temperature. The amount of current they can carry depends on both temperature and the surrounding magnetic field.

The figure here-under shows an example of a cable and busbar configuration.



Figure 8.6. Superconductors Cable Cross-Section
(source OCP: Advanced Datacenter AC/DC Distribution and Conversion Power Architecture)

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

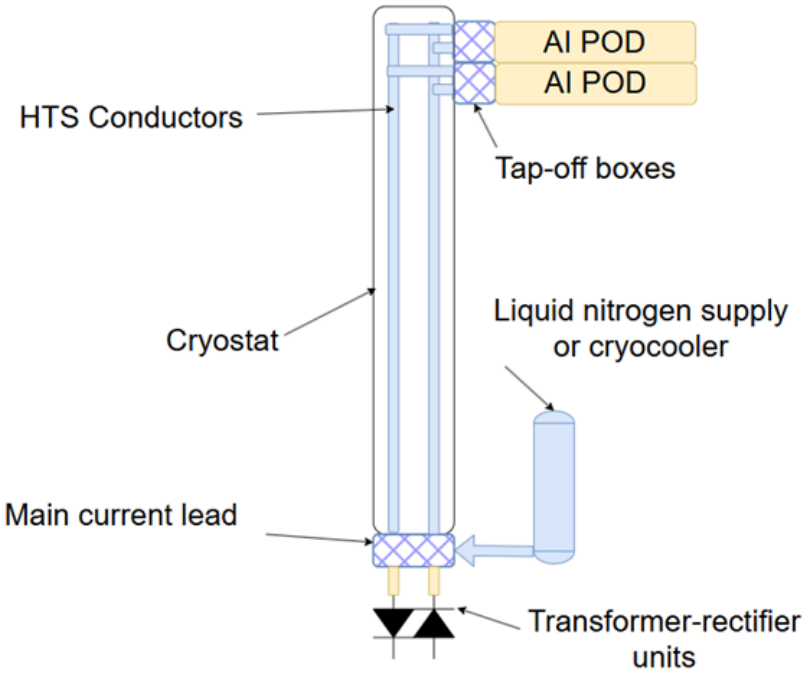


Figure 8.7. Cooling Heat Exchange Loop

Superconductor conductors enable significantly smaller, lighter, and more compact power delivery to the data hall whitespace compared with conventional cables and busway systems. These reductions translate directly to improved overhead congestion management, increased layout flexibility, and support for higher-density compute deployments.

The zero electrical resistance of the superconductors enhances the overall efficiency of the system and reduces heat generation inside the building, hence less stress on the HVAC system. However, it requires a specific dedicated cooling system to cool down the conductor at very low temperatures, which means that this solution would be only economically valuable for very high power and long run of conductors.

Superconductors have not yet been deployed in data centers, but the technology is under active development. The individual components required have already been designed and tested, although not specifically for data center applications. The ongoing trend toward higher rack power densities and DC power architecture strongly favors superconducting solutions. As these trends continue and MW+ racks are built, the deployment of the first superconducting data-center systems will become more and more economical.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

Another promising solution, especially for the connection to the Rack itself would be the use of liquid cooled cable similar to the one used in the EV industry for MCS type of application (Megawatt Charging System). More of this type of cable is presented in the Connectors sections.

The principal applicable standards are:

- UL857 → Busway
- IEC 61439-6 → Busway
- There are standards which cover traditional cables.
- Standards cover superconducting cables:
- IEC 63075 - Superconducting AC power cables and their accessories for rated voltages from 6 kV to 500 kV: test methods and requirements.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

- Connectors



Figure 8.8. Example of an IEC 60309 connector

Currently there is no applicable standard for separable connectors in the range up to 1500 VDC and currents in the 1 kADC range for data center applications. Products may be designed and released following e.g. IEC 60309 / UL 1682, IEC 61984, or UL 1977 (probably in conjunction) with UL 4128 – or other applicable standards – supplemented by application specific risk analyses in order to cover the extended operating range [REF TO CODES AND STANDARDS].

However, such connectors exist in other industries as mentioned above in the conductors section. UL 2278 covers cable and connectors for MCS for the EV industry with rating up to 1500 Vdc and 3000 A, which makes this solution a good candidate to be applied inside an AI Datacenter. EV Charging and AI Datacenter are 2 different applications, which means that to be able to use this type of technology, the industry must ensure the requirements written today for the EV charging application cover those of AI Datacenter.



Figure 8.9. Example MCS Connector

Conditions for Normal Operation and under Failure conditions must be defined. One important Normal Operation Condition to be defined is ability or non-ability for disconnect under load. While disconnect under load is technically feasible via hybrid electronic switching or other means, the economy of such a feature should be challenged because of the resulting size and cost. For disconnect under no-load conditions pilot contacts should be incorporated to support

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

interlock means on the load side (rack) or on the DC grid side (circuit breakers) linked with pilot or control pins or microswitches. This way circuit breakers can shut down strains to ensure safe disconnecting. Also the type of operator should be considered. If there are skilled operators, interlocking mechanisms preventing unwanted disconnection under load might be suitable. If there are unskilled operators, more sophisticated solutions to ensure safe disconnecting under load should be implemented.

For breaking capacity under load it should also be considered if full operational current and voltage need to be disconnected. Often devices can be shut down to idle modes before disconnecting. Therefore a connector rated for example 800V and 125A might only need a breaking capacity of 800V but only 2 to 5A, which is in a range that can be handled using different protection methods.

Important parameters for Operation under Failure Conditions are overload and short circuit conditions, which are imposed on the connector from the embedding system, i.e. the DC grid and the load (racks). The Joule Integral I^2t and I_{peak} which the connector must sustain under short circuit conditions must be determined on system level and covered by constructive means (conductor cross sections, magnetic force capability, etc.).
Consideration for handling capability (cable pulling weight and flexibility, overhead or blind mating, mechanical locking, etc.) should be considered.
Alternate termination means hard wiring with cable lugs, bolts, etc. Such methods, however, might require a certain training and certification level of the operator/installer while a connector typically can be operated with minimum training and allows for more speedy installation.
In data centers high capacitive loads are to be expected. Therefore connecting live parts result in high inrush currents. These peaks might harm contacts or can trigger protection devices like fuses or circuit breakers. When disconnecting and connecting under load is a general use case the effects of voltage peaks as a result of disconnecting and high inrush currents during connecting must be considered on system level.

8.6. AI Rack

For detailed explanation of IT racks refer to Chapter 9.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

8.7. Miscellaneous

In this section are listed equipment which have not been listed as Main Component but play a significant role in the system.

1) Cooling

   CDU (Cooling Distribution Unit) embeds equipment which could run on DC power. In some cases, feeding CDU in DC would result in better overall efficiency, for example when removing the rectifier stage of a Variation Frequency Drive. It is also important to consider redundancy/backup for CDU power given critical loads are just as dependent on their cooling as they are on their direct power for reliable operation and mitigation of downtime.

2) Metering/Monitoring

   As per traditional AC power system, metering and monitoring could be required at different areas (main DC output, feeder, branch). The requirements for such features will be discussed in a future version (and may be incorporated in other OCP component base specifications, such as those for power and BBU shelves).

3) Operation/Maintenance

   Equipment related to operation and maintenance (e.g continuous thermal monitoring) are not covered in this paper version. Readers can refer to NFPA70B for more information.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

9. Interface with Rack & Power

This section provides an overview of a rack architecture supporting LVDC input connection, most of the content is regarding LVDC input components and vertical bus bar, however lower voltages in this type of rack is possible by means of an LVDC input conversion shelf, architecture and design of 50VDC IT trays and bus bar is well defined today in Open Compute so they are only mentioned as options but not extensively explored.

There are three options to connect the LVDC distribution to a LVDC input Rack that provide different flexibility depending on the needs of the end user, all of them coming from a tap box and DC whip combination, however the use of a bus bar with a flexible braided section is an option to support higher currents. Figure 9.1 shows these three options with more details in the following subsections.

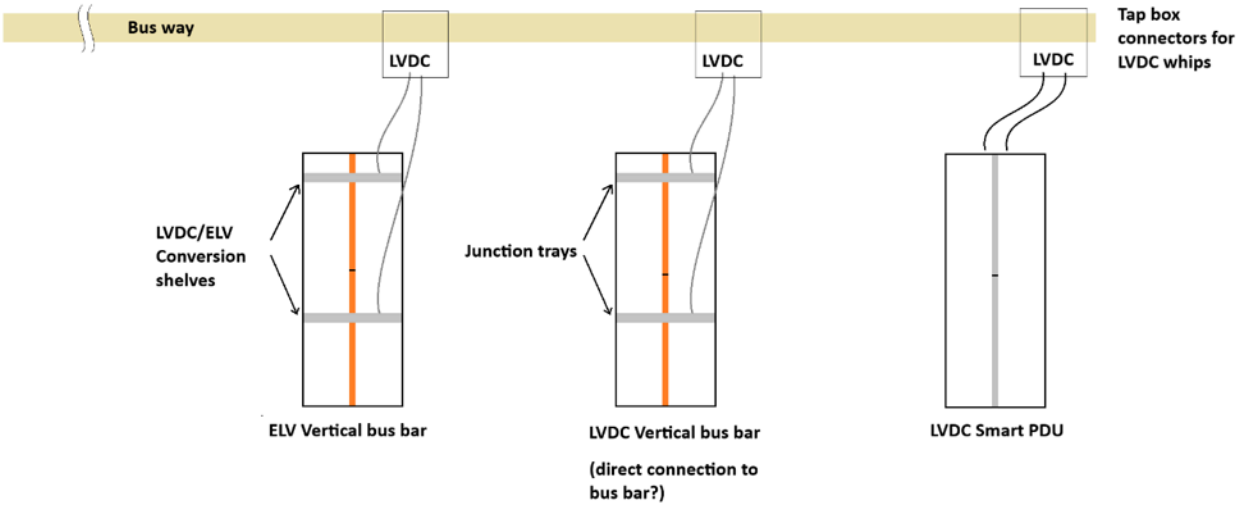


Figure 9.1. LVDC Rack input connection options

Figure 9.1 left shows a connection from the tap box to a LVDC/ELVDC conversion shelf to power conventional 50V input IT trays e.g. compute trays, switch trays, storage trays, etc. The LVDC/ELVDC conversion shelf shall include reinforced insulation from primary to secondary side since typical IT trays using 50V do not incorporate isolated converters. This option provides the capability of utilizing existing IT infrastructure in a Datacenter with LVDC distribution if necessary. The rack can be divided into one or more power domains, each one with its own LVDC/ELVDC conversion shelf, as an example two power domains are shown in this figure.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

Figure 9.1 center shows an example of a LVDC input rack using junction trays to connect from the tap box to the internal vertical bus bar of the rack, supplying LVDC to all the LVDC-native IT trays in the rack; each one of the IT trays needs to have its own LVDC/ELVDC conversion with reinforced insulation and hot swap capability. The rack can be divided into one or more power domains, each one with its own DC whip as input, two power domains are shown in this figure as an example.

Figure 9.1 right shows an example of a rack utilizing a LVDC smart PDU to power LVDC-native IT trays, this option provides the capability of individually controlling turn ON/OFF individual nodes and including features like circuit breaker and monitoring.

9.1. Components of an LVDC Input Rack

9.1.1. Junction Tray/Box

For racks with LVDC vertical bus bar it is recommended to incorporate a junction tray or junction box in the rack to function as input power interface between the tap box in the bus way and the bus bar in the rack, although is possible to incorporate this interface on the rear of the bus bar, having it provides the flexibility to incorporate additional rack level monitoring and protection features.

9.1.2. LVDC Vertical Bus bar

The vertical bus bar installed in the rack needs to be designed based on whether bipolar or unipolar input is being used. In the case of bipolar input, the bus bar shall incorporate four different conductors L+, M, L- and PE that come from the LVDC distribution at datacenter level. In the case of unipolar three conductors shall be used, L+, L- and PE. The bus bar conductors (except for PE) shall be galvanically isolated from the rack and shall incorporate touch-safe features to prevent accidental contact.

9.1.3. LVDC Smart Rack PDU

Instead of a vertical bus bar an LVDC Smart Rack Power Distribution Unit (PDU) can be installed on the sides of the rack to provide power to LVDC-native IT trays, these type of PDU can incorporate the typical features of a Smart Rack PDU, for example a power switch per

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

outlet for remote ON/OFF operation, current, voltage and power monitoring per outlet and for the entire PDU, data logging, etc.

9.1.4. LVDC Input Conversion Shelf

An option for the rack is to incorporate an LVDC input Conversion Shelf to convert from LVDC down to ELVDC (e.g. 50VDC) or even from LVDC to LVDC, to decouple the requirements of the LVDC grid from the ones in the rack, providing several benefits and the ability to scale the LVDC distribution at datacenter level to larger voltages, the potential combinations for input-output voltages of this type of conversion shelf are listed below.

- ±400VDC/800VDC to 50V (input from LVDC distribution).
- ±400VDC/800VDC to 50V (input from side power rack e.g. Mt. Diablo).
- ±700 VDC to 50V.
- ±700 VDC to ±400VDC/800VDC.
- AC input to ±400VDC/800VDC (rectifier power shelves to retrofit in an existing datacenter with AC distribution).

For example, a conversion shelf of this type can incorporate enough bulk capacitors to support transients like voltage dropouts of longer duration and wider input voltages, while providing a tighter regulation at its output and grounding system translation (e.g. from IT to TN-S) which allows to optimize the power density in LVDC input PSUs for LVDC-native IT trays, or keep using compute trays with previous power architectures using ELVDC voltages.
Different models of conversion shelves can be designed to operate with different input voltage ranges e.g. the ranges described in Chapter 5, hence operating as a translation layer between the different LVDC standards, moreover, features like input current limiting facilitate power excursion mitigation when combined with in-rack Capacitor Bank Units (CBUs).

Given the relatively larger volume compared with an LVDC input PSU in an LVDC-native IT tray, the LVDC input conversion shelf provides more space for EMI input filters and can incorporate larger input protection devices e.g. TVS or Varistors while at the same time incorporating fast-acting output over current protection mechanisms, reducing the transients that can be induced in the LVDC distribution at datacenter level.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

The potential tradeoffs of this would be that some space in the rack would need to be used and the conversion shelf represents another conversion stage with losses, however the industry can focus on optimizing the efficiency of such conversion shelf.

9.1.5. Rack Architecture With an LVDC to ELVDC Conversion Shelf

There is no clear definition of when a rack needs to transition from lower voltage (e.g. 50VDC) distribution in its vertical bus bar to a higher voltage (e.g. LVDC), and the selection of ELVDC or LVDC as vertical power distribution in the rack is done based on the needs of the specific designs, however the more power needed in the rack the higher will be the current that needs to be supported by the vertical bus bar and the amount of copper that needs to be used to handle the current. Higher current bus bars in ELVDC rack distribution can incorporate liquid cooling methods to mitigate the temperatures generated in the bus bar due to its losses. The figure 9.2 below provides an example of popularity in the industry depending on the power level of the rack.

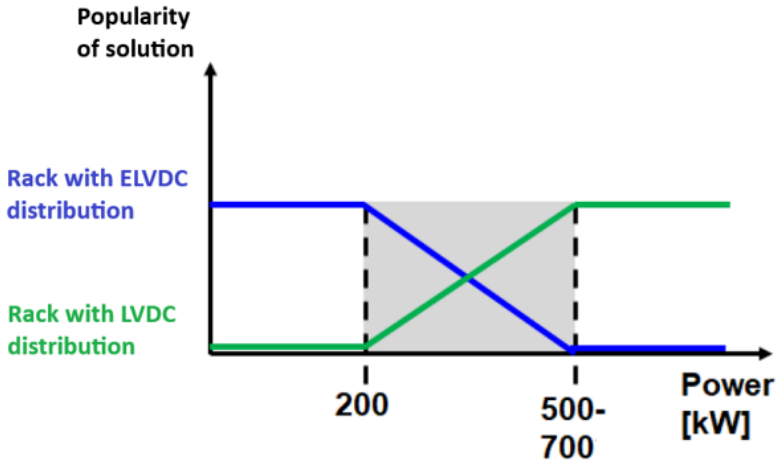


Figure 9.2 Example of popularity of a rack with ELVDC vertical bus bar (e.g. 50V) and a rack with LVDC vertical bus bar for power distribution

Rack power levels for LVDC vertical distribution can include 200-250 kW, 500 kW, 750 kW and potentially more in the future.

Using LVDC vertical distribution in a rack introduces new requirements and areas to take care of, the following subsections provide an overview of such requirements.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

9.2. Rack Architecture with an LVDC to ELVDC Conversion Shelf

Figure 9.3 shows the power architecture of an IT rack with LVDC to ELVDC conversion/power shelf. In this architecture, the IT rack is supplied by a low-voltage DC (LVDC) input. The LVDC input is converted to extra-low-voltage (ELV) DC by the LVDC to ELVDC power shelf. It should be noted that in an IT rack, multiple DC-DC power shelves can be connected in parallel as required to support the load and/or redundancy. The ELVDC at the output of LVDC to ELVDC power shelf is then delivered to the server/compute tray via a vertical busbar. The server/compute tray includes an e-fuse/hot-swap circuit followed by an onboard DC-DC converter. This onboard converter steps the ELVDC down to an intermediate bus voltage (IBV) DC level. Finally, the IBV DC is converted to a sub-1 V level by the voltage regulator module (VRM), which supplies power to the loads e.g. CPU, GPU, NICs, Storage, etc.

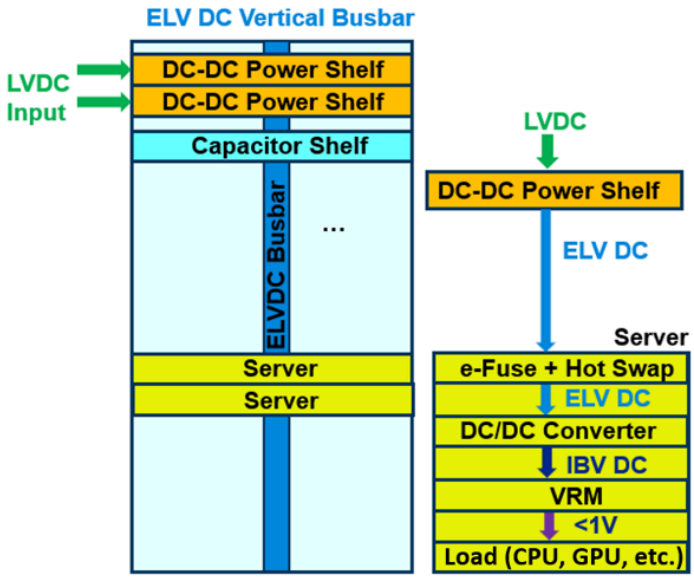


Figure 9.3 Power architecture of IT rack with LVDC-ELVDC power shelf

The LVDC to ELVDC power shelf consists of multiple paralleled, galvanically isolated DC-DC power supply units (PSUs). Each DC-DC PSU includes its own e-fuse and inrush-current-limiting device. In the event of a fault in a specific DC-DC PSU, the e-fuse requires several milliseconds to isolate the fault, while the inrush-current-limiting device helps limit the fault/inrush current and prevents an excessive drop in the LVDC bus voltage.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

Additionally, to prevent a significant LVDC bus voltage sag during the several-millisecond interval before the e-fuse opens and isolates the faulty DC-DC PSU, additional hold-up capacitors are placed at the LVDC input. As a result, the LVDC to ELVDC power shelf must satisfy a hold-up-time requirement in the several-millisecond range.

To maintain a stable ELVDC bus voltage and prevent fast GPU load transients from propagating upstream and causing oscillations in the LVDC and AC input of upstream AC-DC converter (SST), a capacitor shelf can be placed on the ELVDC busbar. The capacitor shelf placed on the ELVDC busbar provides energy buffering and peak load shaving for fast GPU load transients.

9.3. IT Rack With LVDC-native IT Tray

Figure 9.4 shows the power architecture of an IT rack with an internal power supply (PSU) or power distribution board (PDB) located on the server tray. In this architecture, the LVDC input is distributed directly to the server trays through an LVDC busbar inside the IT rack. Each server tray includes a fuse and hot-swap circuit that provides inrush-current control during start-up, discharges the PSU/PDB input capacitors during shutdown or decommissioning, and enables fault monitoring, fault isolation, and telemetry. It should be noted that in case of any fault, the hot swap circuit on the server tray typically isolates the fault within a few microseconds (µs).

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

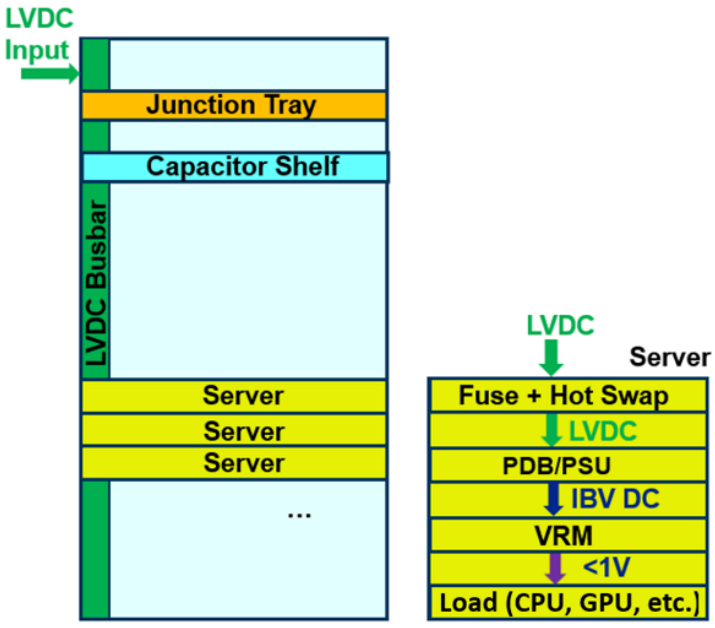


Figure 9.4 Power architecture of IT rack with onboard PSU/PDB on IT tray

Downstream of the fuse and hot-swap circuit, the LVDC input is directly converted to an intermediate bus voltage (IBV) DC level by a galvanically isolated DC-DC converter in the PDB. Finally, the IBV DC is converted to a sub-1 V level by the voltage regulator module (VRM), which supplies power to the GPU load.

Additionally, the IT rack shown in the previous figure includes a junction tray that houses capacitors used to stabilize the LVDC input voltage to the rack. These capacitors help damp oscillations caused by the parasitic line inductance of the LVDC busway.

Furthermore, the IT rack can include a capacitor bank shelf (CBU) which helps to stabilize LVDC bus voltage and prevent fast GPU load transients from propagating upstream and causing oscillations at the AC input of upstream AC-DC converter (SST). The capacitor shelf placed on the LVDC busbar provides energy buffering and peak load shaving for fast GPU load transients.

9.4. Preliminary Rack Level Voltage Specifications

The purpose of this section is to provide an overview of the specifications needed in the rack including preliminary voltage ranges.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

9.4.1. Input Voltage Range

The figure 9.5 below shows the proposal of the voltage ranges for the operation of the components inside the rack using a ±400VDC/800VDC voltage level. These components can be power supplies inside IT trays, capacitor bank units (high frequency and low frequency) as well as any other component attached to an LVDC bus bar in the rack.

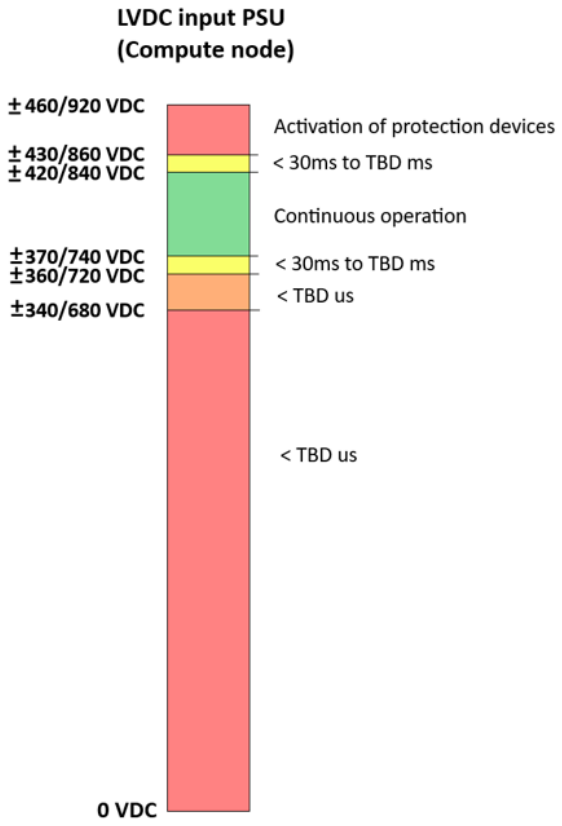


Figure 9.5 ±400VDC/800VDC Input voltage range proposal for a LVDC Input Component (PSU, Conversion shelf, etc.)

9.4.2. Continuous Operation Definition

Continuous operation is defined as a voltage range in which the converters inside the LVDC/ELVDC conversion shelf can operate continuously without tripping any protection e.g. overcurrent or overtemperature.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

9.4.3. Transient Level 1 Operation Definition

PSU or LVDC/ELVDC conversion shelf can operate momentarily without tripping any protection e.g. overcurrent or overtemperature, while providing enough voltage for the load to keep operating normally without performance degradation (throttling). It is expected for this type of transient to last for a few tens of milli-seconds and are related to the opening of fuses on the LVDC distribution at datacenter level.

9.4.4. Transient Level 2 Operation Definition

Transient 2 (orange) operation is defined as the voltage range in which the LVDC Input PSU or LVDC/ELVDC conversion shelf can operate during a few hundreds of micro-second range without tripping any protection e.g. overcurrent or undervoltage while providing enough voltage for the load to keep operating normally without performance degradation. Operation during this type of transient is intended to ride-through events like short circuits inside a nearby compute rack.

9.4.5. Transient Level 3 Operation Definition

Transient 3 (red) Surge is defined as the voltage range in which the LVDC input PSU or LVDC/ELVDC conversion shelf can operate during tens micro-second range without tripping any protection e.g. overvoltage, while providing a proper voltage the load to keep operating normally without performance degradation, activation of protection devices is expected in this voltage rage to clamp the voltage and prevent damage to the LVDC input PSU or conversion shelf.

Transient 3 (red) Dropout is defined as the voltage threshold below the LVDC Input PSU shall have enough energy to provide hold-up time, keeping its outputs with enough voltage for the load to keep operating normally, performance degradation during this condition is allowed to maximize the hold-up time and ride through this condition. There can be a difference between the LVDC input PSU and the LVDC/ELVDC conversion shelf regarding the total energy they can support hence duration of dropout condition given the volume available, having the option of supporting longer dropout times in a LVDC/ELVDC conversion shelf compared to a LVDC input PSU due to its larger volume.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

9.5. LVDC Input Power Supply Architecture and Components

9.5.1. Overview

LVDC-native systems incorporate an internal power supply unit (PSU) what itself is comprised of different components e.g. EMI input filters, Hot Swap Controller circuitry, decoupling and/or bulk capacitors, LVDC/ELVDC conversion stages with different output voltages as required, and a standby converter, the generalized architecture of this type of PSU is shown in the figure 9.6 below.

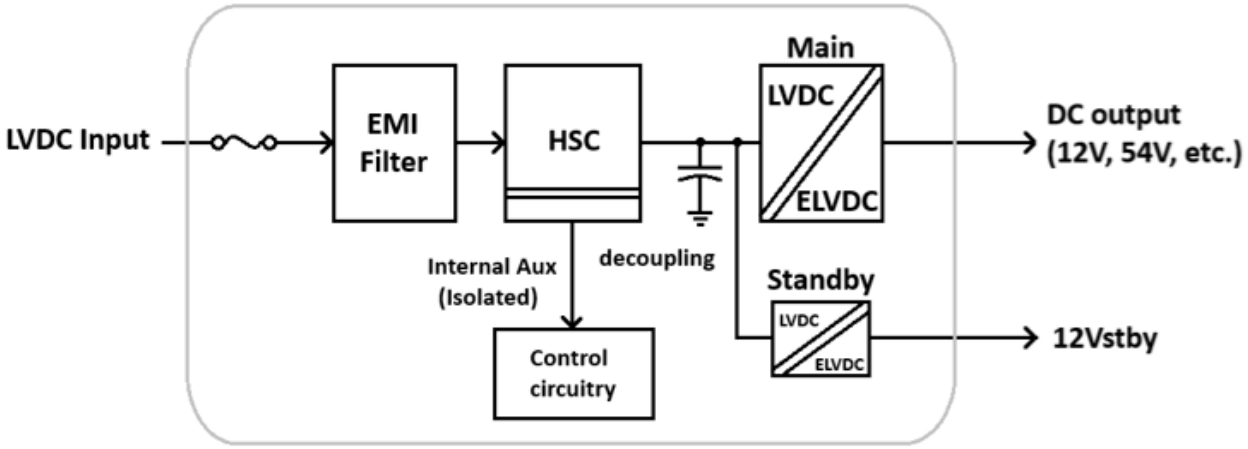


Figure 9.6 Generalized LVDC input PSU architecture

Note 1: Other features like TVS, internal Microcontroller, current monitoring circuitry, etc. are not shown in the figure for simplicity.

Note 2: EMI input filter can be at the input of the individual converters instead of a single input filter, the performance needs to be evaluated by the PSU implementer.

9.5.2. IT Tray Input Connector

The input connector of an LVDC-native IT tray shall incorporate 3 groups of contacts: power contacts, pre-char contacts and presence contacts. Power contacts are the ones carrying current to and from the IT tray e.g. L+, L- and M or function as protection e.g. Protective Earth (PE or Ground), in a bipolar system. Pre-charge contacts are longer contacts that help to charge the parasitic capacitances in the wires and other components. Presence wires act as presence detection to determine when the IT tray is fully inserted or is being extracted so the HSC

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

circuitry can quickly turn OFF its output/outputs and void having a disconnection under load, these pins can also be used for input voltage sensing, necessary for telemetry in the IT tray. The Figure 9.7 below shows an example of a bipolar and unipolar bus bar clip type connector with all the described groups of contacts and their insertion sequence.

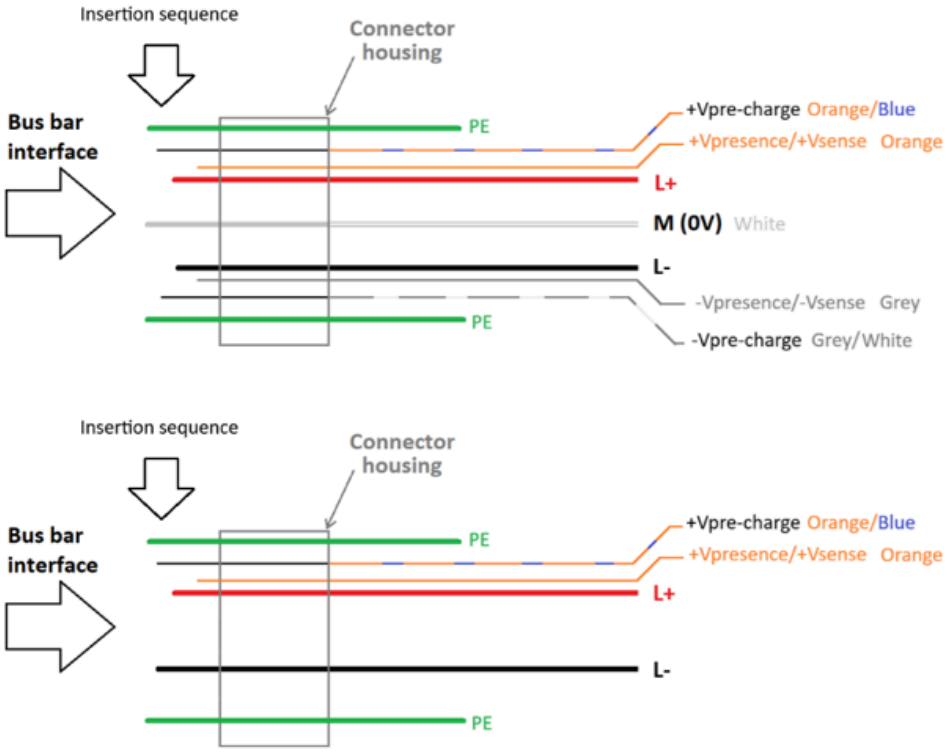


Figure 9.7 Bipolar (top) and unipolar (bottom) bus bar clip connectors for IT trays

9.5.3. Hot Swap Controller Circuitry

The Hot Swap Controller (HSC) is the circuitry that allows any type IT tray to be inserted into the rack while the bus bar is energized without creating a high inrush current that can be generated when trying to charge the internal decoupling/bulk capacitors. The HSC provides a controlled current source to charge the decoupling/bulk capacitors present in the LVDC input PSU. It also provides fast acting over current protection necessary to quickly turn OFF the internal DC bus in case of a short circuit failure on any of the components connected to it, one example can be a short circuit in the input of a LVDC/ELVDC converter module.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

The figure 9.8 below shows a conceptual block diagram of an HSC circuitry incorporating the pass-through elements (transistors), shunt resistors to measure the current, the control circuitry and the isolation barrier needed to communicate with the ELVDC side,

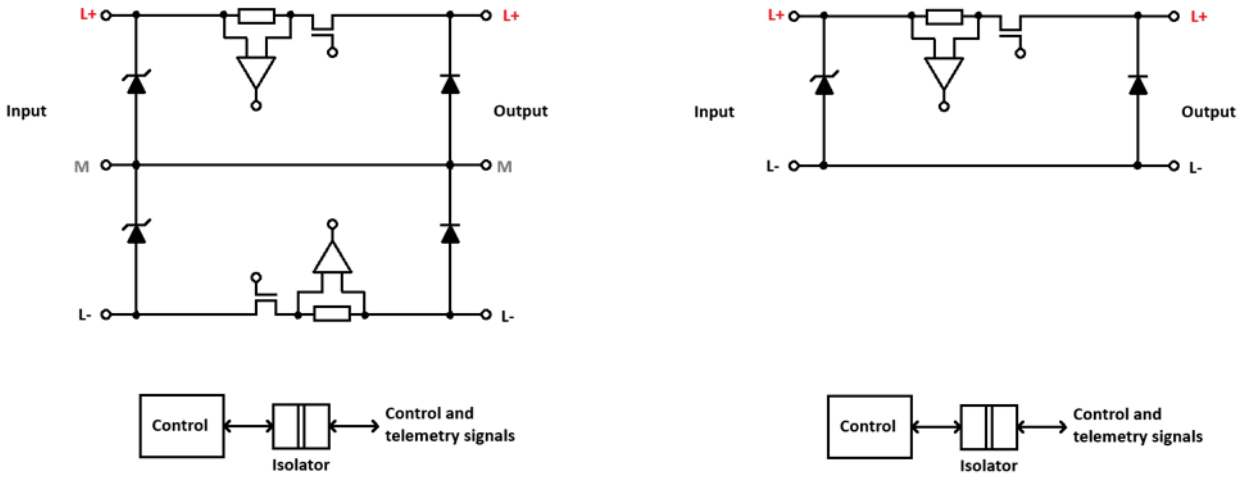


Figure 9.8 Conceptual HSC block diagram, bipolar (left), unipolar (right), aux. power supply not shown

The functions and features described along with a small auxiliary power supply can be incorporated into a Hot Swap Controller module, keeping the complexity in it, moreover a modular design provides reusability for different types of IT trays.

9.5.3.1. Initial Inrush

independently of having an HSC circuit at the input and no EMI input filter circuitry, this is because the parasitic capacitances in the wires, PCBs, transistors, protection devices, etc. usually this type of initial inrush current events last for about 1us to 3 us, however the current can be significant often surpassing 30 Amps.

The current and timing of this event will depend on the design of the LVDC input PSU and the wiring inside the IT tray, to help to mitigate this, it is recommended to have a parasitic capacitance charger or pre-charge circuit that is comprised of a longer pin the input connector along with a current limiting device (e.g. resistor) that can limit the current to a controlled level,

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

the evaluation of such type of pre-charge circuit needs to be done by the implementer of the IT tray. The figure 9.9 below shows an example of initial inrush current without pre-charging.

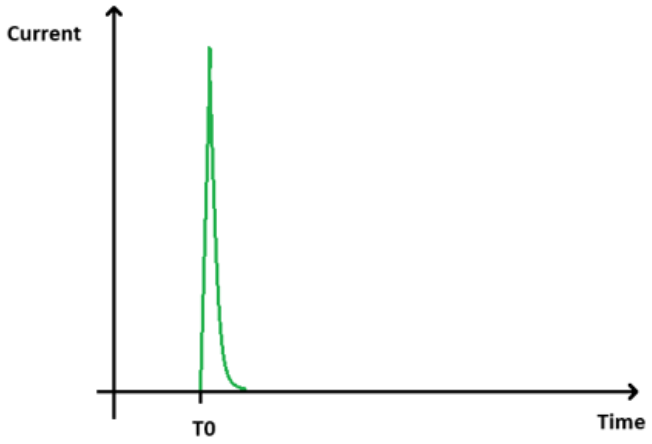


Figure 9.9 Example of an initial inrush current without pre-charging

9.5.3.2. Charging Current and Startup Sequence

Figure 9.10 below shows an example of how the input current evolves during the different stages, initially when the IT tray is inserted represented by T0, the input current will show an initial inrush limited by the pre-charged circuit to an I2 current level defined by the implementer, once the parasitic capacitances have been charged, the input current will go down to close to zero current, some current will be still be present due to leakage of some components, resistive voltage dividers in the HSC circuity and as the major contributor the input current of an auxiliary power supply (e.g. flyback) to bias the HSC circuitry.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.



Figure 9.10 Input current during different stages

decoupling/bulk capacitors inside the LVDC input PSU, for the established duration of time, it is important that this charging current is kept relatively low to reduce the stress in the pass-through devices. This current will depend on the capability of the devices in the HSC circuitry however it is recommended to keep it below the nominal operating int current of the IT tray.

After the decoupling/bulk capacitors have been charged at T2, there will be an additional input current at the input, this increase is due to the auxiliary power supply that LVDC/ELVDC converters need to operate.

When the LVDC input PSU is enabled (turned ON) and load is applied at T3, the input current will increase to a level that depends on the load, and it is represented by I3 in the figure.

9.5.3.3. Monitoring and Telemetry

It is recommended that the HSC circuitry incorporate monitoring and fault detection with features like: Overvoltage protection (OVP), Undervoltage Protection (UVP), Overcurrent protection (OCP), Short circuit protection (SCP) and Overtemperature protection (OTP) as well as telemetry features to provide the information to the IT tray.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

9.5.3.4. Discharge

LVDC input PSU shall incorporate a discharge circuit to discharge the internal decoupling/bulk capacitors down to a safe voltage in a time specified by the safety standards.

9.5.3.5. Internal Auxiliary Power Supply

Control and sensing feature that this might have e.g. Microcontrollers and amplifiers, it is recommended to bias this auxiliary converter from the main input before the HSC circuitry such that, in case of a failure, this auxiliary converter can keep providing power to the devices in the HSC circuitry as well as the internal Microcontroller.

9.5.4. DC/DC Converters (LVDC/ELVDC)

LVDC input PSU inside an IT tray incorporates DC/DC isolated converters (LVDC/ELVDC) that provide the conversion down to the voltages required for the load e.g. 50V, 12V, etc. These converters will have to be sized according to the load needs depending on the IT tray implementation.

9.5.4.1. Connection Recommendations

LVDC/ELVDC converters can be connected using different configurations, depending on the input voltage of each one of them and the type of regulation they use e.g. fixed-ratio or regulated output, figure 9.11 below shows the different connections for a bipolar LVDC input PSU.

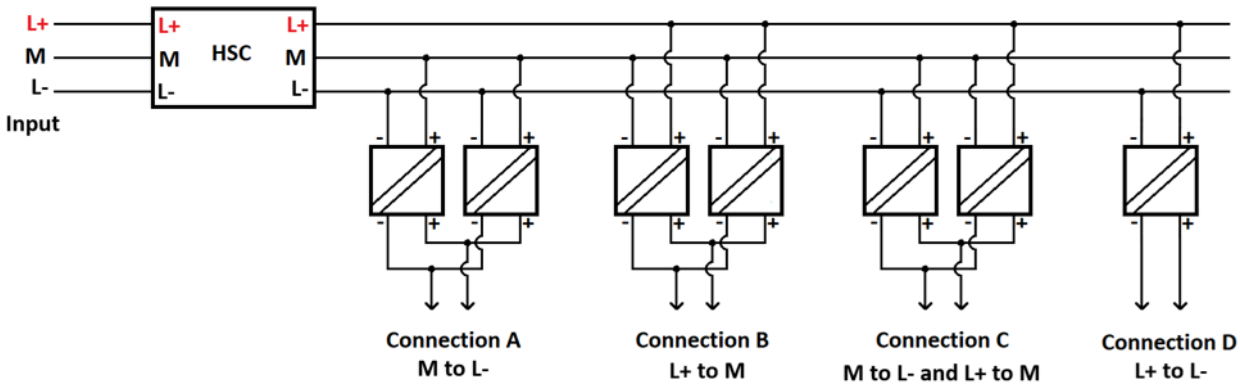


Figure 9.11 Examples of converter connection options using a bipolar LVDC input PSU

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

The connections A & B are recommended for fixed ratio converters, these connections provide the modules with the same input voltage, reducing the impact on the output current share error. Connection C is recommended for regulated output converters, fixed-ratio converters will not work using this connection. Using regulated output converters as in connection C will help to balance the currents in the L+ and L- conductor, however there will be some current through the M wire given the current share error of the converters. Connection D can be used if a unipolar input converter is used in the LVDC input PSU however when this connection is used, the input fuses (not shown) at the input of the HSC as well as the transistors need to be rated for the L+ L- potential e.g. 800VDC when ±400VDC is used.

When using a unipolar LVDC input PSU the connection can be done as shown in the figure 9.12 below.

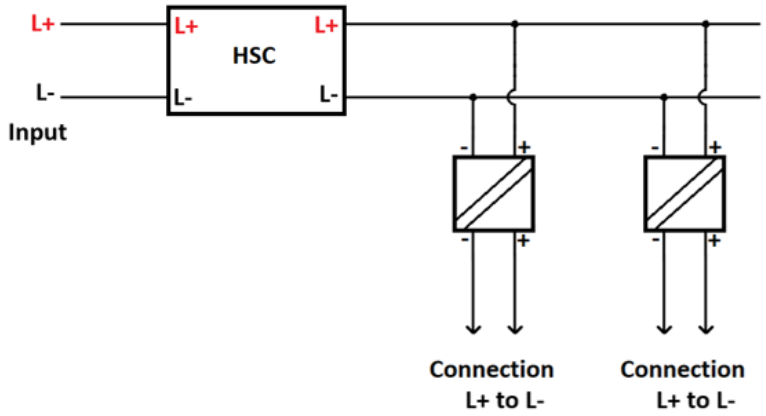


Figure 9.12 converter connection using a unipolar LVDC input PSU

9.5.5. Isolation

Galvanic reinforced isolation is required for the LVDC/ELVDC converters from primary to secondary, this includes the main converters and any other auxiliary converter, as well as any telemetry, sensing or control feature that communicates the information from the LVDC side to the ELVDC side and vice versa.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

9.5.6. IT Tray Grounding

The LVDC input power supply shall incorporate a Protective Earth (PE a.k.a. Ground) connection in the case, sized according to the applicable safety standards. Figure 9.13 below shows an example of the ground connection inside the IT tray and the LVDC input PSU as well.

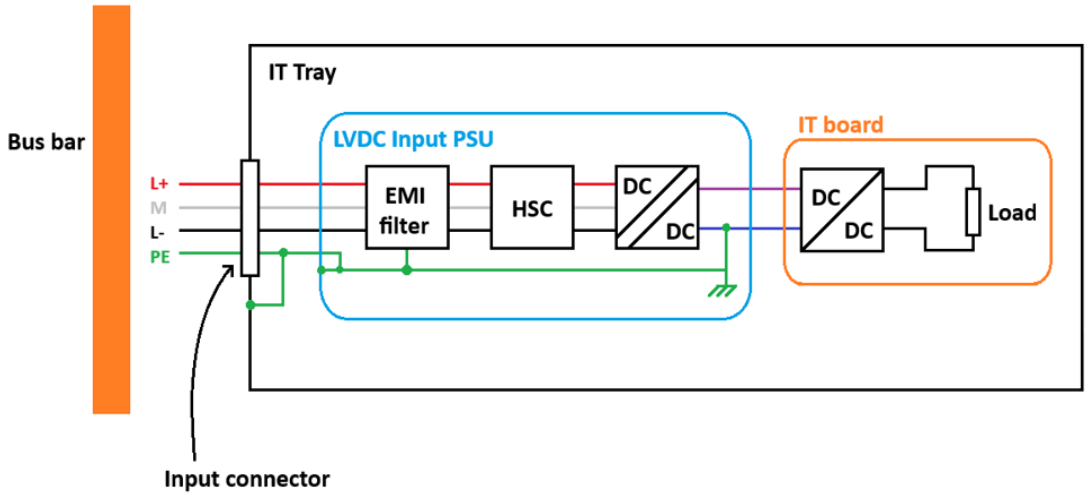


Figure 9.13 Example grounding system in an IT tray and a bipolar LVDC input PSU

Note: for Unipolar LVDC input PSU the M wire is not present.

9.6. Energy Storage and Power Oscillation Mitigation

The rapid adoption of AI and accelerated computing is fundamentally changing datacenter power system requirements. Unlike conventional IT loads, modern GPU clusters exhibit highly dynamic power profiles characterized by large and frequent load transients. The power consumption of AI processors can vary from near-idle to full-load conditions within milliseconds, generating significant high-frequency power fluctuations and large di/dt stress on the DC distribution infrastructure.

As datacenter power levels continue to scale from tens of megawatts to hundreds of megawatts, these load fluctuations can propagate through the DC power distribution network and create undesirable effects, including:

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

- DC bus voltage oscillations
- Increased current stress on cables and busbars
- Reduced power quality at the rack and row level
- Interaction with converter control loops
- Grid-side power oscillation through SST or AC/DC front-end systems
- Reduced system stability and reliability
- Oversized (capacity) grid-side BESS

To maintain power quality and ensure stable operation, energy storage resources must be strategically deployed throughout the datacenter power architecture. These storage resources serve different functions depending on their location within the power hierarchy and their response speed requirements.

A multi-layer energy storage architecture is therefore required, ranging from ultra-fast storage located close to the GPU load to larger energy storage systems located at the distribution or utility interface level.

9.6.1. Rack Level Energy Storage To Handle The Pulsating Power

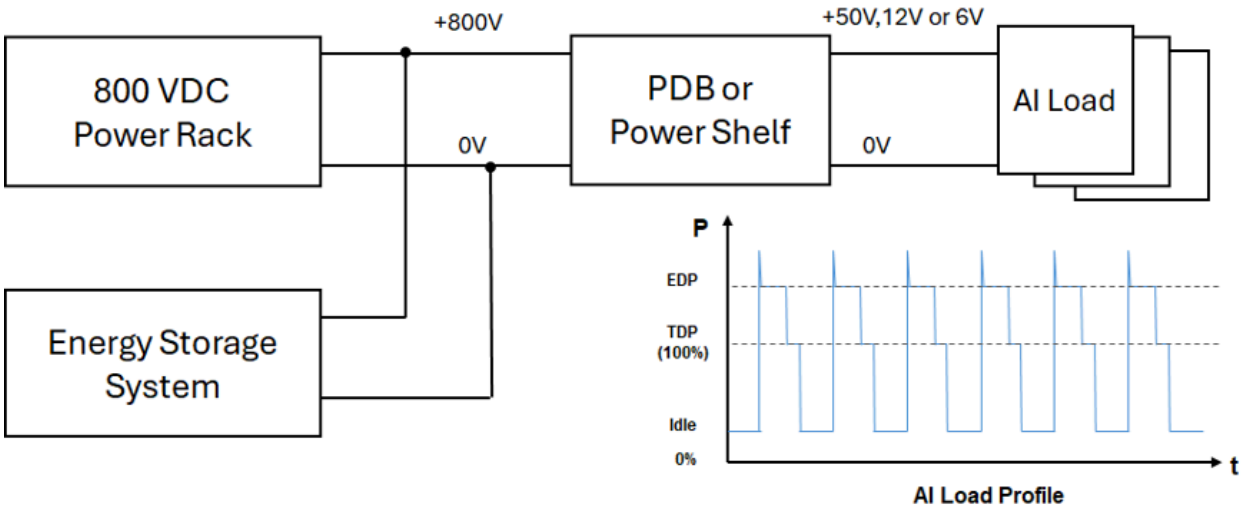


Figure 9.14 Rack Level Energy Storage to handle the quick di/dt

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

The fastest power fluctuations originate directly from GPU and AI accelerator workloads. These power variations typically occur in the millisecond range and contain significant high-frequency components.

Allowing these high-frequency power oscillations to propagate throughout the DC distribution system can result in:

- Excessive bus voltage ripple
- Higher RMS current in busbars
- Increased losses in power conversion equipment
- Larger converter current ratings
- Potential interaction among multiple racks operating simultaneously

To mitigate these effects, fast-response energy storage should be deployed as close as possible to the load. At the rack level as shown in Figure 9.14, the energy storage system absorbs or delivers power during rapid load transitions, effectively decoupling the AI load from the upstream power infrastructure. During a sudden GPU power increase:

- The local storage system immediately provides transient energy.
- The upstream power shelf supplies only the average power demand.
- DC bus voltage excursions are minimized.

During a rapid reduction in GPU power:

- Excess energy is temporarily absorbed by local storage.
- Reverse power flow and voltage overshoot are reduced.

This approach significantly reduces the dynamic current demand seen by the rack power distribution system.

Several technologies can be considered for rack-level deployment:

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

Technology	Response Speed	Energy Density	Power Density	Typical Application
Capacitor	Microseconds	Low	Very High	High-frequency ripple suppression
Supercapacitor	Microseconds to milliseconds	Medium	High	GPU load transient support
Battery	Milliseconds to seconds	High	Moderate	Short-duration ride-through
Battery Capacitor	Microseconds to minutes	High	High	Integrated ride-through and load transient support

For AI datacenters, local storage prevents unnecessary propagation of transient power disturbances throughout the datacenter electrical infrastructure.

9.6.2. Distribution Level Energy Storage To Meet The Grid Requirements

While rack-level storage addresses fast transient behavior, larger-scale energy storage is required to manage lower-frequency power fluctuations and support grid-facing requirements. As datacenter power capacities approach utility-scale levels, operators increasingly face requirements related to:

- Grid stability
- Frequency regulation
- Voltage ride-through
- Power ramp-rate control
- Renewable energy integration
- Emergency energy backup support
- Utility interconnection compliance

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

Some of the detailed ESS functions are discussed in Section 8.4 of this report.

9.6.3. Hierarchical Energy Storage Architecture

A future AI datacenter is expected to adopt a hierarchical energy storage architecture, where different technologies operate at different timescales.

Storage Layer	Response Time	Primary Function
Capacitor Tray	µs	Voltage ripple suppression
Supercapacitor Tray* *functionality may be absorbed into row-/distribution-level storage	ms	GPU transient compensation
Row-Level Storage	10 ms–1 s	Cluster power balancing
Distribution Storage	1 s–minutes	Grid support and ride-through
Utility-Scale BESS	minutes–hours	Energy management and backup

This layered energy storage approach minimizes total storage cost while maximizing overall system performance, power quality, and grid compatibility. Energy storage can be strategically deployed at multiple levels within the power architecture, including the computing tray, computing row, beginning of DC distribution interface, and even at the medium-voltage (MV) AC grid connection point. Each storage layer serves a distinct function, enabling optimized responses across different time scales and power dynamics.

The integration of Solid-State Transformers (SSTs), DC distribution, and hierarchical energy storage creates scalable and future-proof architecture for next-generation AI datacenters. By combining fast-response local storage with larger centralized energy resources, the system can effectively manage highly dynamic computational loads while maintaining stable operation and minimizing stress on the utility grid.

This architecture provides a practical pathway toward higher energy efficiency, improved grid interoperability, enhanced resiliency, and reduced infrastructure costs. Furthermore, it enables

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

the evolution from traditional passive power delivery systems to intelligent, grid-interactive datacenters capable of supporting hundreds of megawatts of rapidly varying AI workloads. Through coordinated control of power conversion, distribution, and energy storage assets, future datacenters can become active participants in grid services while simultaneously delivering the performance, reliability, and scalability required by advanced AI computing infrastructure.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

10. System Operation, Control and Monitoring

10.1. MOP and SOO in Traditional AC Data Center Architectures

In traditional AC data centers, operational safety and reliability depend on two distinct paradigms: MOP and SOO.

MOP (Method of Procedure): A MOP is a step-by-step, rigidly documented manual procedure executed by human operators. In AC systems, human-driven MOPs are highly effective because AC power (with its natural zero-crossings) and traditional equipment (with large thermal and mechanical inertia) provide wide tolerance windows—often ranging from seconds to minutes. Operators rely on MOPs for macro-level system reconfigurations, scheduled maintenance, and physical Lock-Out/Tag-Out (LOTO) protocols, ensuring safety through strict human verification.

SOO (Sequence of Operation): Conversely, an SOO is the automated, hard-coded control logic executed by equipment-level controllers (like PLCs or DSPs). It governs actions that are too complex, fast, or critical for human intervention. In AC environments, SOOs manage processes requiring precise synchronization or sequential staging, such as UPS phase-locking during bypass transfers, generator paralleling, or the step-by-step startup of chiller plants.
In summary, MOPs dictate the macro-level human-to-system interactions across flexible timeframes, while SOOs govern the micro-level machine-to-machine reflexes requiring strict millisecond-to-second precision.

10.2. Transitioning to the DC Paradigm: From Human-centric MOP to System-centric SOO

As data centers evolve towards 800V DC architectures, the balance between Method of Procedure (MOP) and Sequence of Operation (SOO) shifts fundamentally. In traditional AC systems, the MOP is the primary driver, relying on human operators to manage equipment with significant physical inertia. However, in the LVDC AI datacenter world with added buffering capacitance, human reaction time is an order of magnitude too slow to prevent semiconductor device damage or manage arc-less switching.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

Consequently, SOO is expanding while MOP is diminishing. Most critical operational sequences—such as millisecond-level pre-charging, active discharge, and ultra-fast fault isolation—are being hard-coded into the firmware of power converters. The operator's role (MOP) is transitioning from "hands-on execution" to "high-level oversight," where the SOO autonomously handles the micro-level physics of the 800V bus to ensure system survival.

10.3. The "Wild West" of LVDC Native Datacenter Control: Undefined Boundaries

While SOO replaces MOP in 800V DC data centers, the industry faces a critical challenge: the exact boundaries and implementations of DC SOOs remain dangerously undefined.

Undefined Boundaries: There is no consensus on where local equipment firmware (Edge SOO) ends and central management (DCIM SOO) begins. For instance, should fault recovery be orchestrated centrally or handled autonomously by local micro-controllers?

Fragmented Implementations: Without universal standards, vendors rely on proprietary, black-box firmware for microsecond-level reactions (e.g., private communication protocol for coordination vs. local triggers).

The Architect's Challenge: The hardest part isn't deploying HVDC hardware; it's writing the missing rulebook to ensure multi-vendor interoperability without sacrificing the microsecond reflexes essential for system survival.

10.4. SOO Scope for an MVAC-to-LVDC Data Center

An MVAC-to-LVDC data center requires the SOO to expand from equipment-level start/stop logic into a deterministic control and operations rulebook. The SOO assigns ownership for every fast energy event, stored-energy transition, and permissive that could expose personnel or equipment to an unsafe 800 VDC condition, defining which layer acts autonomously, which supervises, which records and alarms, and which actions remain under operator MOP control. This division is essential because the consequential DC events occur in microseconds to seconds, whereas human MOP execution is appropriate only for authorization, verification, maintenance boundaries, and recovery oversight.

The SOO should treat the MVAC-to-LVDC power train as an integrated system spanning the utility point of common coupling, MV switchgear, transformer, solid-state transformer (SST), active front end (AFE) or rectifier, DC bus, bus ties, DC protection devices, BESS or other

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

energy storage if used, row/rack power distribution, cooling interlocks, EPMS/DCIM coordination, and the operator interface. Each element carries a defined state machine and exchanges only the commands and statuses required for safe coordination. Black-box vendor firmware may still perform local microsecond protection, but the site-level SOO should specify the observable states, required inputs, permissives, inhibits, event logs, reset rules, and fail-safe behavior expected of that firmware.

SOO Vision: SOO is a first-order design artifact for an LVDC data center and is developed in parallel with the one-line diagram, protection study, grounding study, small-signal / impedance stability study, control-network architecture, utility-interconnection design, and commissioning plan. The critical outcome is a deterministic rulebook that preserves the microsecond reflexes required for DC system survival while giving the owner, operator, and utility a transparent operating framework for safety, reliability, maintenance, and recovery.

10.5. Architectural Control Boundaries & Interoperability

10.5.1. Control boundaries

Referring to the table, the control boundary can be organized as a layered autonomy model. Semiconductor and gate-drive protection remains local and deterministic. Equipment controllers manage converter pre-charge, current limiting, contactor/breaker coordination, and internal lockout. A bus or zone controller coordinates bus sectionalization, selective isolation, staged load admission, and local recovery. A site controller or EMS coordinates MV source availability, energy-storage dispatch, redundancy, capacity, and load-shed policy. DCIM and the operator authorize operating modes, maintenance windows, capacity policies, and return-to-service decisions. Any function required to prevent equipment damage, personnel hazard, or uncontrolled DC bus collapse should not depend on cloud services, DCIM availability, non-deterministic communications, or human reaction time.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

Table 10.1 Control boundaries

Layer	Primary scope	Autonomy requirement	Typical response domain
L0 semiconductor / gate drive	Desaturation, shoot-through, overcurrent, overvoltage, thermal trip, gate block	Must act locally without network or supervisory approval	Nanoseconds to microseconds
L1 equipment controller	AFE, rectifier, SST, DC/DC converter, BMS, DC breaker, rack power shelf	Must protect itself and place equipment in a safe state on local evidence	Microseconds to milliseconds
L2 bus / zone controller	Pre-charge, bus-tie logic, sectionalization, selective isolation, rack admission, local recovery	Must coordinate zone survival and isolate faults even if EPMS/DCIM is unavailable	Milliseconds to seconds
L3 plant controller / EPMS	MV source selection, generator/BESS dispatch, redundancy state, capacity, load shedding	May optimize and coordinate, but is not relied upon for immediate fault survival	Seconds to minutes
L4 DCIM / operator MOP	Work authorization, maintenance windows, capacity policy, alarm response, return-to-service approval	Provides oversight and authorization rather than fast protective action	Minutes to hours

10.5.2. Interoperability and Communications Protocols

The MVAC-to-800 VDC data center SOO should define an interoperability architecture for all protection, control, monitoring, and supervisory communications between MV switchgear, protective relays, rectifiers, DC/DC converters, 800 VDC bus controllers, DC breakers, energy storage systems, rack power systems, EPMS/BMS/DCIM platforms, and the utility interface.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

Standardized interfaces should be required wherever practical to avoid proprietary black-box coordination and to support multi-vendor interoperability, deterministic event analysis, cybersecurity, commissioning, and lifecycle maintenance. Minimum Interoperability requirements should include:

- Common point map and naming convention to avoid vendor-specific ambiguity in alarms, states, trips, and permissives
- Common state model - Ensures “ready,” “blocked,” “faulted,” “discharged,” and “maintenance safe” mean the same thing across vendors
- Protocol ownership matrix - Defines which protocol is used for protection, control, monitoring, diagnostics, and utility interface
- Latency and determinism requirements - Prevents using slow supervisory links for fast protective functions
- Time synchronization standard - Enables sequence-of-events reconstruction after DC faults
- Fallback behavior on communications loss - Prevents unsafe operation during network, controller, or gateway failures

10.5.3. Protocol Domains

Domain	Likely protocol / interface class	SOO role
MV switchgear, relays, protection IEDs	IEC 61850, GOOSE, MMS, sampled values where applicable	Fast protection signaling, breaker status, interlocks, event records
Utility / substation SCADA	DNP3 / IEEE 1815, IEC 61850, utility-specified interface	PCC telemetry, utility ride-through, trip/reclose, demand response
Utility DERMS / grid-edge / demand response	IEEE 2030.5, DNP3, utility-specified interface	Utility operating limits, demand response, load curtailment, ride-through coordination, DER dispatch
Plant controllers, PLCs,	OPC UA, Modbus RTU/TCP,	Supervisory control,

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

Domain	Likely protocol / interface class	SOO role
EPMS/BMS	BACnet/IP, vendor APIs	monitoring, alarms, setpoints
MV AC Input Rectifier and other Rectifiers, DC/DC for Battery Back Up, BESS PCS, SSCBs.	Digital Interface Protocol for Electronic Instrument Transformer, Real-time Industrial Ethernet Fieldbus Protocol, CAN2.0B and vendor defined proprietary protocols	System state machine, fault management, Interlock, Operation and maintenance, System re-energization
Rack power shelves / IT load controls	Redfish, SNMP, Modbus, vendor APIs	Rack admission, power capping, staged load shed
Time synchronization	IEEE 1588 PTP / IEC 61850-9-3 where required	Time-aligned sequence-of-events and fault analysis

10.5.4. Communications, Cybersecurity, and Event Evidence

The SOO needs to separate protection-critical signaling from supervisory monitoring. Local protection remains functional during DCIM loss, network outage, cloud outage, or supervisory-controller reboot. Loss of communications causes each layer to enter a defined conservative mode rather than an undefined fallback. Commands are rejected if stale, unauthorized, duplicated, out of sequence, inconsistent with local status, or issued during an active inhibit. Time synchronization is required for event analysis, with high-resolution event records for pre-charge attempts, contactor and breaker operations, converter trips, insulation alarms, energy storage transitions, load-shed events, ride-through events, and discharge verification.

Cybersecurity is treated as part of operational safety, because an unauthorized close, blocked trip, false discharge status, or manipulated protection setting can create a physical hazard. The SOO defines command authority, role-based access, dual confirmation for high-risk commands, signed-firmware expectations, version compatibility, protection-setting change control, audit trails, and a loss-of-trust mode. Multi-vendor interoperability includes not only data mapping but common definitions of safe state, lockout, reset, inhibit, and verified discharge.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

10.6. Deterministic Sequence of Operations (SOO)

10.6.1. System Operation, Control and Monitoring Aspects

The main aspects to be developed and hard-coded into the operational sequence include:

- Sequence Of Operation
- Startup Sequence (Black Start, Precharging)
- Hot Connection/Disconnection (discharging)
- Energy Measurement
- Power management / voltage “droop curves”

10.6.2. MVAC Source Acceptance and LVDC Bus Energization

The SOO defines the full sequence from MVAC source acceptance through stable LVDC bus availability. Source acceptance verifies MV voltage magnitude, frequency, phase rotation, phase balance, relay health, breaker status, grounding status, transformer or SST readiness, communications status, and absence of downstream DC faults. Closing the MV breaker or enabling the rectifier is inhibited by an active work permit, LOTO condition, closed ground switch, active discharge path, unresolved DC fault, failed insulation monitor, failed voltage sensor, or unavailable protection device. Transformer energization accounts for inrush behavior, differential protection status, thermal permissives, and upstream utility switching limitations. An SST uses pre-charge, intermediate-link control, output regulation, bypass, inhibit, discharge, and lockout sequences across its AC/DC and DC/DC stages.

10.6.3. Startup Sequence

10.6.4. Black Start

[Content to be expanded based on specific site architecture.]

10.6.5. Precharging

A DC grid contains many capacitances; when uncharged capacitors are connected to a DC voltage, very high inrush currents flow which would cause the protective devices to switch-off without further measures. It therefore is essential to ensure that the inrush current is limited to levels that will not activate protective devices at power-up. Capacitor voltage needs to be ramped up gradually by a precharging current that does not exceed the rated current of the respective sector being switched-on.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

Preferably, a precharging device measures voltage at its input terminals and output terminals, respectively. When the voltage at the output terminal of the devices is within a tolerance band around the operational voltage, precharging is complete and the sector is turned on.
During load transfers from one bus to another, it will be important to have the busses at the same voltage level to avoid high peak currents and oscillations.

The LVDC bus pre-charge is a necessary SOO sequence. Before pre-charge, the system confirms that downstream rack feeds are open or soft-start blocked, the bus is not tagged for maintenance, no temporary grounding or discharge device is connected, and insulation monitoring is healthy. The SOO then enables the pre-charge path, ramps bus voltage under current limit, validates the expected voltage-rise profile, verifies polarity and voltage-sensor agreement, closes the main DC contactor or breaker only after voltage differential is within limits, removes or bypasses the pre-charge element, and declares the bus available. A failed pre-charge produces specific diagnostic outcomes, including suspected shorted bus, open pre-charge element, welded contactor, reversed polarity, failed voltage measurement, high leakage, or unexpected downstream capacitance.

10.6.6. Power Management with Droop Curves

In direct current systems a disbalance in power supply and power demand leads to a change of voltage that is measurable at any point in the installation. Higher demand than supply leads to lower voltage and vice-versa. Since DC systems do not have to stick to a fixed voltage value (this is a significant advantage over the rigid frequency demands of AC systems) this behavior is used to balance the power demand again.

The means to do so is active components such as DC/DC converters measuring the voltage at their terminals and adjusting their current output according to defined current-voltage curves, known as “droop curves”. These curves, in the most basic and most robust variant, are pre-defined and are immune to outside interference. Besides regulating the power supply, also loads can be adjusted. Preferable those with a larger time constant such as heating or cooling devices.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

10.6.7. Rack Load Admission, IT Coordination, BESS, and Cooling Interlocks

High-density AI racks are admitted to the 800 VDC bus as controlled electrical loads rather than uncontrolled capacitance and step load. The SOO defines rack-admission permissives including available bus capacity, source redundancy, cooling availability, rack power-shelf health, insulation status, and absence of local fault. Racks are admitted in stages or waves to limit aggregate inrush and dI/dt. The system supports workload-aware power capping, priority-based load shedding, and brownout modes that preserve controls, networking, storage, and management functions before large GPU load blocks. After a grid or bus disturbance, rack re-admission is delayed and staggered to avoid a second bus collapse or a utility ramp-rate violation.

Where an energy storage device is present, the SOO should define whether they support the 800 VDC bus, the MVAC point of common coupling, or both. Storage is inhibited from feeding a faulted or maintenance-tagged bus. Energy storage dispatch observes state of charge, thermal limits, charge/discharge limits, fault-current contribution limits, utility import/export restrictions, and recharge ramp limits. A battery thermal alarm, BMS fault, fire/life-safety interface trip, or loss of required ventilation or cooling isolates the energy storage and blocks automatic restart until the required inspection or reset condition is met.

Cooling is an electrical permissive in an 800 VDC AI data center. A rack or row is not energized if the required CDU, pump, flow, valve state, supply temperature, leak detection, or heat-rejection path is unavailable. Cooling degradation triggers power capping or staged load shed before equipment reaches thermal trip. A leak detected in a rack, row, or CDU zone triggers the defined electrical isolation sequence for the affected equipment while preserving healthy zones where the protection design allows.

10.6.8. Utility Interface, Ride-Through, and Grid-Disturbance SOO

Ride-through needs to consider converter-control, DC-bus-stability, protection, and utility-interconnection. Where power conversion, energy storage, and generation resources operate in parallel with the utility, the SOO defines the applicable ride-through basis before assigning setpoints. IEEE 1547 applies when interconnected with the electric power system at the distribution interface, while IEEE 2800 applies where resources are interconnected with transmission or sub-transmission systems. For IEEE 1547, the SOO identifies the selected

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

abnormal-performance category, voltage and frequency trip settings, voltage and frequency ride-through curves, rate-of-change-of-frequency behavior, enter-service and reconnection requirements, automatic-reclosing coordination, and any mutually agreed dynamic-voltage-support functions. For IEEE 2800, the SOO defines voltage and frequency ride-through, high- and low-voltage disturbance behavior, dynamic voltage support, dynamic active-power support during abnormal frequency, reactive-power or real-power priority during current limiting, negative-sequence response during unbalanced conditions, phase-angle-jump ride-through, RoCoF ride-through, consecutive voltage-dip performance, momentary-cessation or current-blocking limits, post-disturbance power recovery, and return-to-service behavior. The SOO distinguishes events where the plant stays connected, events where it may ride through using energy storage or controlled load reduction, and events where it intentionally separates to protect the facility or comply with the interconnection agreement.

The ride-through sequence classifies the grid condition as normal, disturbed but ride-through capable, disturbed requiring support mode, disturbed requiring load shed, grid lost, grid return pending, resynchronization or retransfer, and post-disturbance recovery. During a shallow sag, the rectifier or SST may continue operating with controlled current limiting and DC-bus voltage regulation. During a deeper sag or mandatory ride-through event, energy storage may support the 800 VDC bus while nonessential rack admission is blocked and noncritical load is reduced according to the load-priority table. During a phase-angle jump, RoCoF event, phase-unbalance event, or unstable utility recovery, the SOO prevents nuisance trip, loss of converter synchronism, uncontrolled inrush, and simultaneous restart of large rectifier blocks. During utility reclose or source transfer, the SOO staggers converter re-enablement, blocks rack re-admission until the MV source is stable, and confirms that protection, synchronization, harmonic, and power-factor permissives remain satisfied. In large AI data centers, sudden loss or return of load can itself become a utility event, so the SOO coordinates import ramp rate, BESS discharge and recharge limits, harmonic limits, reactive-power or power-factor mode, voltage-support mode, and any utility-required load-shed or load-restoration envelope.

Where connected sources can operate in parallel with the utility, the SOO should address anti-islanding, backfeed prevention, synchronization, reconnection timers, voltage and frequency ride-through settings, RoCoF and phase-angle-jump response, current-blocking or momentary-cessation limits, and utility protection interfaces. It is reconciled with the

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

interconnection agreement, the applicable IEEE 1547 or IEEE 2800 category or performance requirement, the adopted utility ride-through curves, harmonic and flicker requirements, protective relay settings, and any utility-specific voltage, frequency, reclose, and restoration requirements.

10.6.9. DC Bus Sectionalization, Protection, and Fault Isolation

Each bus-tie close requires voltage-magnitude agreement, polarity agreement, acceptable voltage differential, compatible grounding condition, no active fault, no active maintenance tag, and sufficient short-circuit capacity. Bus-tie opening is defined for fault isolation, maintenance isolation, source reconfiguration, and capacity management. The SOO also prevents unintended paralleling of incompatible sources, incompatible converter control modes, or independently grounded sections. In this context, “incompatible control modes” is a small-signal stability condition, not only a settings mismatch: tightly regulated rack and load converters can present a negative incremental input impedance (behaving as constant-power loads) over certain frequency bands, and off-the-shelf converters that are each stable in isolation can interact and cause instability once interconnected on a common bus, because the equivalent source and load impedances depend on how many units, of what topology and control loop, are paralleled. To keep paralleling decisions deterministic in a multi-vendor environment, the recommended pathway is to treat a converter-level impedance/passivity margin as an interconnection acceptance criterion. For example, requiring a positive real part of the converter terminal impedance across the frequency range of interest, or an equivalent source-to-load impedance (minor-loop-gain) margin. In this way, the bus-tie and source-paralleling permissives can be evaluated against a defined stability boundary rather than left to vendor defaults.

DC fault isolation is a coordinated sequence. For a pole-to-pole short, local converter current limiting and gate block act first, while the nearest appropriate DC breaker, hybrid breaker, solid-state breaker, or fuse isolates the faulted section before upstream trips remove healthy capacity. For a pole-to-ground fault, the response follows the selected grounding philosophy, defining first-fault alarm, second-fault trip, high-resistance-ground current threshold, or solid-ground clearing behavior as applicable. For an arc fault, the SOO isolates the affected section, suppresses upstream energy contribution, blocks automatic reclose, and requires inspection before reset. After any repeated fault or protection-miscoordination event, automatic recovery is inhibited and elevated to operator review.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

The SOO defines the relationship between converter protection and discrete DC protective devices. Converter gate blocking stops semiconductor contribution rapidly but may not remove all stored or backfed energy; a breaker or fuse can isolate a section but may not act quickly enough to prevent local semiconductor damage. The SOO therefore coordinates current limits, crowbar or dump paths (if used), contactor sequencing, breaker trip commands, fuse selectivity, arc detection, upstream source blocking, and post-trip discharge.

10.6.10. Thresholds and Failure Timing Budget

The SOO is deterministic only when its permissives, alarms, and trips are expressed as quantified setpoints. The values below are representative reference figures consistent with current OCP 800 VDC and LVDC power-distribution work (including the Diablo rack and power base specification) and with applicable safety and power-quality standards. Final setpoints are intended to be fixed by the project-specific protection, short-circuit, and grounding studies and reconciled with the utility interconnection agreement. The below table establishes the parameter set the SOO defines, with typical magnitudes and the governing reference.

Unlike AC systems, an 800 VDC bus has no current zero-crossing and very low inductance, so fault current can reach destructive levels in microseconds and the bus can collapse before a conventional, time-graded scheme would clear. Classic AC selectivity, which tolerates 100–200 ms coordination delays, therefore cannot be transplanted. The SOO must instead achieve selectivity primarily through current-limiting and threshold/zone discrimination at electronic speed, with mechanical and time-graded devices providing slower backup. The timing budget below allocates, end to end, who acts and by when for a pole-to-pole fault on a distribution bus.

10.7. Grounding, Insulation Monitoring, and Personnel Safety Interfaces

800 VDC grounding philosophy should be stated explicitly in the SOO, because ground-fault behavior depends on whether the system is floating, high-resistance grounded, solidly grounded, midpoint grounded, bipolar (±400 VDC), or unipolar (800 VDC). The control system performs insulation checks before energization, continuously monitors insulation during operation, alarms or trips per the selected philosophy, and supports ground-fault location where practical. A first ground fault in an IT-style floating system may be alarm-only while a second fault requires immediate isolation; a solidly grounded system may require faster clearing

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

because ground-fault current can be higher. These distinctions are captured in the SOO rather than left to equipment defaults.

Personnel safety interfaces include emergency stop, remote trip, access-door interlocks where used, maintenance grounding status, verified absence of voltage, discharge confirmation, and LOTO permissive logic. A device does not declare maintenance safe based on an open command alone; it requires source isolation, downstream backfeed prevention, discharge verification, no rebound, and the local or independent verification required by the site safety program. The SOO also should define how shock and arc-energy warning states are presented to operators, so that alarm language is unambiguous.

10.8. Stored Energy, Active Discharge, and Verified-Zero-Energy

The SOO should include a formal stored-energy management sequence for all 800 VDC bus sections, converter DC links, energy storage interfaces, and rack-level capacitance. A controlled shutdown blocks new rack admission, reduces controllable IT load where possible, opens downstream feeders in the prescribed order, disables converter gating, opens source-side devices, and then enables the active discharge path. The discharge sequence monitors voltage decay against an expected profile, detects voltage rebound, supervises discharge-device temperature, and confirms voltage is below the defined safe threshold before declaring verified discharge. An abnormal decay profile triggers an alarm, maintains lockout, and requires investigation rather than return to service.

The MOP boundary requires independent verification before physical work but does not replace the SOO discharge logic. The operator receives explicit equipment states such as discharging, discharge failed, verified discharged, and maintenance safe. Return to service is equally controlled: remove temporary grounds, clear work permits, verify all isolation devices and access panels, confirm no active inhibits remain, re-run insulation and voltage checks, perform pre-charge, and re-admit load in a staged sequence.

10.9. Minimum SOO Deliverables

The expanded control and operations package is delivered as a set of authored SOO documents rather than a single narrative. At minimum it includes SOOs for MVAC source

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

acceptance; transformer, SST, or rectifier energization; 800 VDC bus pre-charge; bus-tie close and open; rack and row load admission; normal shutdown; emergency shutdown; active discharge and verified-zero-energy; pole-to-pole fault isolation; pole-to-ground fault response; arc fault response; converter module failure; energy storage ride-through and isolation if applicable; utility disturbance ride-through; MV source loss; cooling-failure electrical response; communications loss; cybersecurity loss-of-trust; maintenance isolation; return to service; black start; full plant restoration; firmware or settings change management; and commissioning proof testing.

The SOO package is validated before field energization through protection-coordination studies, DC short-circuit studies, grounding studies, harmonic and power-quality studies, small-signal / impedance-based stability studies (source-versus-load impedance and passivity margins, verified by impedance measurement where practical), utility-interconnection requirements, hardware-in-the-loop or real-time simulation where practical, factory acceptance testing, site acceptance testing, and periodic proof testing. The objective is to prove not only that each device protects itself, but that the integrated MVAC-to-800 VDC system transitions safely between states, isolates faults selectively, rides through credible utility disturbances, remains small-signal stable across credible source and load combinations, discharges stored energy predictably, and presents operators with unambiguous states for MOP execution

10.10. Enabling Technologies for Microsecond Survival-New Equipments

10.10.1. DC Fault Transient Recorder

In AC systems, fault recorders typically log waveform changes in milliseconds across electrical cycles. Because Solid State Circuit Breakers (SSCBs) clear faults in a matter of 10 to 25 microseconds, a traditional recorder would entirely miss the event.

- Core Function: Capture ultra-high-frequency transient processes to reconstruct the exact electromagnetic signature of DC short circuits, ground faults, or arc flashes, providing "black box" data for root cause analysis and protection tuning.
- Key Characteristics:
    - **MHz-Level High-Frequency Sampling & Edge Computing:** To keep up with SSCB clearing speeds, recorders will require sampling rates of at MHz. Handling this volume of data necessitates embedded FPGAs for real-time feature extraction.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

The device will process the raw data locally and upload only the critical high-frequency signatures and inflection points to prevent network congestion.

- **Microsecond Time Synchronization & Deep Coupling:** Requires extreme clock synchronization precision (e.g., IEEE 1588 PTPv2). Furthermore, it must be deeply coupled at the hardware level with the SSCB's underlying controller architecture. The moment the SiC MOSFETs or custom IGCTs detect an anomaly, the recorder must instantly lock and buffer the pre- and post-fault waveforms.
- **Wideband Current Sensor Integration:** Standard Current Transformers (CTs) cannot measure DC, and conventional Hall effect sensors lack the necessary response time. This new recorder must interface directly with high-frequency Rogowski coils or high-bandwidth shunts

10.10.2. DC Sequence Controller (SOO Automation Unit)

As DC rectifiers scale up to 1.6MW and beyond, energizing and de-energizing the power train is no longer a simple "close the breaker" operation; it is a highly volatile energy balancing act.

- **Core Function**: Orchestrate the complex startup, pre-charging, state-machine transitions, and post-fault network reconfiguration of the DC bus, ensuring smooth energy transitions and equipment safety during operational shifts.
- **Key Characteristics**:
  - **Rigorous Pre-Charge Timing Coordination:** High-Performance Computing (HPC) nodes are backed...

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

11. Electromagnetic Compatibility (EMC)

This subsection examines the EMC challenges that LVDC presents as they are becoming an alternative to traditional AC architectures for higher power density data centers.

Ensuring EMC compliance is essential not only for regulatory adherence but also for maintaining system reliability and interoperability.

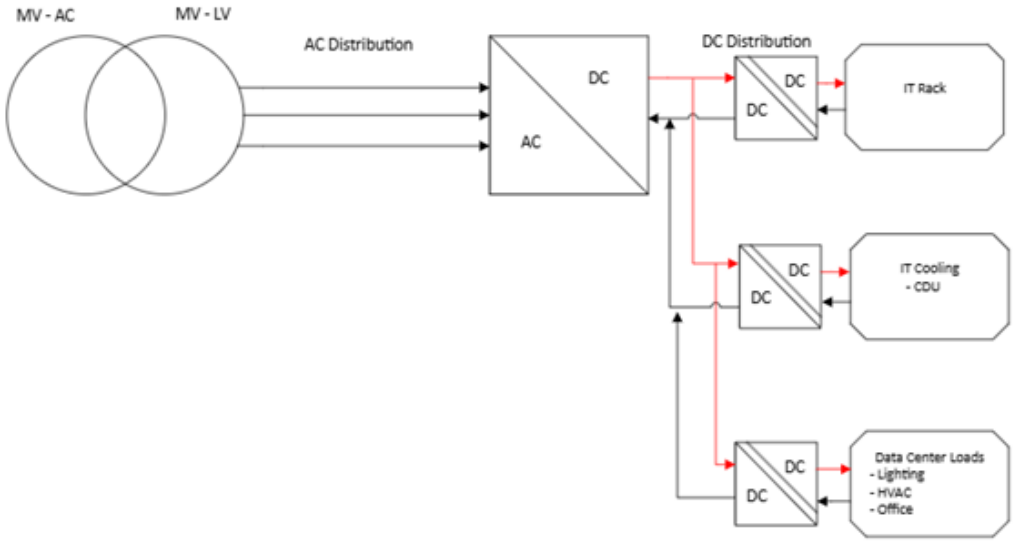
Both AC-DC converters and DC-DC power supplies must meet emissions and immunity requirements, while operating under realistic loading conditions representative of data center environments.

11.1. Scope and System Context

EMC considerations apply across multiple levels of the data center power architecture consisting of:

- AC/DC Converters
- DC/DC Converters
- AC/DC/DC Converters

Other components such as cooling distribution units are outside the scope of this work, but EMC considerations should be considered.



This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

11.1.1. AC-DC Power Converters

Utility interface (including medium voltage systems and solid-state transformers)

Also refer to System Architecture 1.1 MV-LV drawings System Architecture workstream drawings

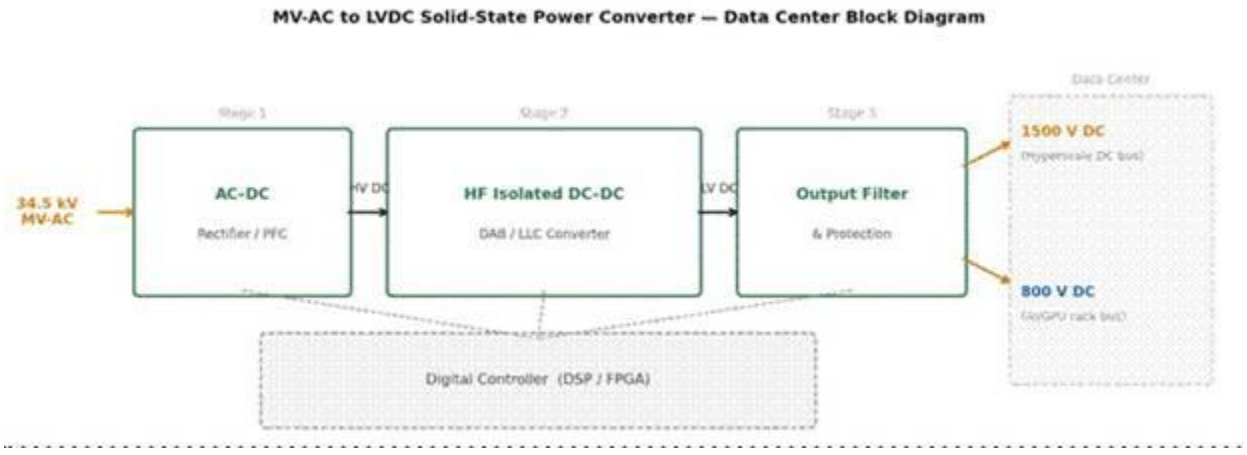
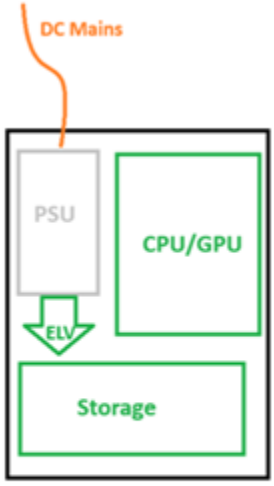


Figure 11.1 34.5 kV MV-AC to 800 V / 1500 V LVDC Solid-State Power Converter

11.1.2. DC-DC Converters

Rack-level and board-level power delivery

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.



Compute node
(top view)

LVDC input PSU developed as a separated unit to be individually certified

11.2. Emissions Requirements

11.2.1. Applicable Standards

- CISPR 11 (industrial environments)
- CISPR 32 / FCC Part 15 (IT equipment)
- IEC 61000-6-4 (generic industrial emissions)
- IEC 61000-3-12 / 3-11 (AC harmonic and flicker limits) - Grey area now however for power supplies between 16-75A this is still an applicable requirement.

*For systems greater than 75A IEC/TR 61000-3-6 can be considered

Data center equipment typically operates under Class A limits, reflecting controlled industrial environments.

11.2.2. Conducted Emissions

AC Interfaces typically from 150kHz to 30MHz. As standard committees modify Class A limits for equipment starting at 9kHz this may need to be revisited.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

Areas that need to be considered for higher voltage/current testing are finding or designing EMI filters and Line Impedance Stabilization Networks (LISNs) to meet the product needs.

DC Interfaces unlike AC systems are not yet fully regulated and represent a key gap in existing EMC standards. Areas to consider are Noise propagation across shared DC buses and Cross-coupling between racks and converts. To help reduce these concerns limits need to be established for DC conducted emissions (e.g., 9 kHz to 30 MHz), and during design validation bus-level noise interactions should be evaluated.

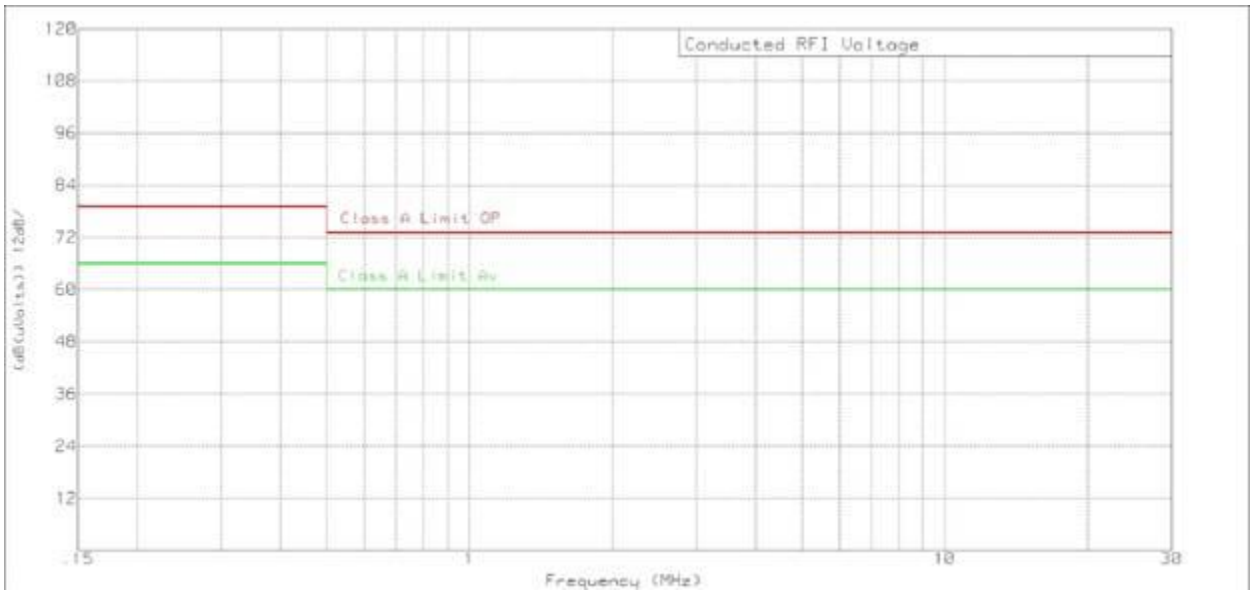


Figure 11.2 FCC Class A and CISPR 11 Class A Group 1 Limits (0.15-30MHz)

11.2.3. Radiated Emissions

Radiated emissions are typically from 30MHz and until at least 1000MHz, but can be up to 40GHz, and are influenced by:

- Switching frequency content
- Package parasitics
- Cable harnesses acting as antennas

Some mitigation strategies for reducing Radiated Emissions include:

- Shielded enclosures

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

- Controlled PCB layouts (loop minimization)
- Proper cable shielding and bonding of the shield

Note: Radiated emissions are primarily driven by switching frequency harmonics, cable harness antenna effects, and enclosure seam resonances. Mitigation through shielded enclosures, controlled PCB layout, and cable bonding is expected.

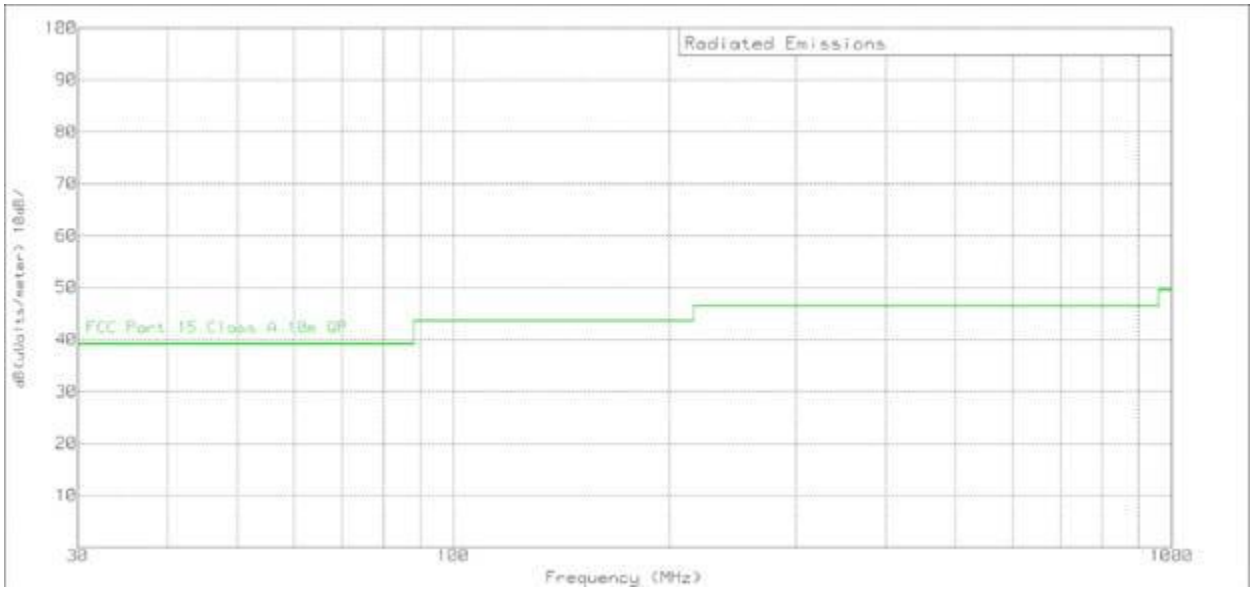


Figure 11.3 FCC Class A Limit (30-1000MHz)

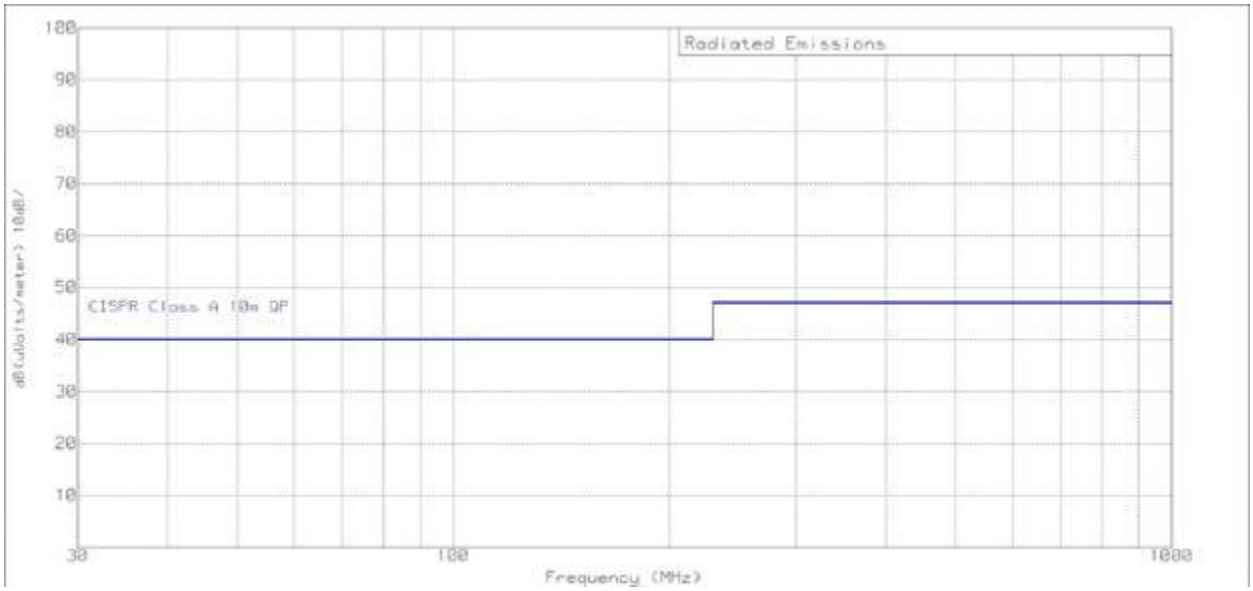


Figure 11.4 CISPR 11 Class A Group 1 (30-1000MHz)

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.



Figure 11.5 FCC Class A Limit (1-18GHz)



Figure 11.6 FCC Class A Limit (18-40GHz)

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

11.2.4. Harmonics and Ripple

Harmonic limits apply primarily to AC inputs, but DC systems introduce unregulated ripple and harmonic behavior, and there is currently no consensus standard for DC ripple limits. To help with this it is considered best practice to limit DC ripple to <1-2%, and validate the impact on downstream converters and controlling loop stability.

11.3. Immunity Requirements

11.3.1. Applicable Standards

- IEC 61000-6-2 / 6-5 (industrial immunity)
- IEC 62040-2 (UPS Immunity)
- CISPR 35 (IT equipment immunity)
- IEC 61000-4-x test methods

11.3.2. Key Immunity Phenomena

Key tests include:

- Electrostatic discharge (ESD)
- Radiated RF fields
- Electrical fast transients (EFT)
- Conducted RF disturbances
- Surge immunity
- Voltage dips and interruptions

Testing must include both AC and DC ports, and I/O or Signal lines as applicable.

Test	Standard	Test Level	Performance Criteria	Limitations
Electrostatic Discharge	61000-4-2	4	B	
Radiated Immunity	61000-4-3	3 or 4*	A	*Higher test levels may be needed based on the installation location

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

Test	Standard	Test Level	Performance Criteria	Limitations
EFT	61000-4-4	4	A	CDN for voltage/current
Surge	61000-4-5	4	B	CDN for voltage/current
Conducted Immunity	61000-4-6	4	A	CDN for voltage/current
Power Frequency Magnetics	61000-4-8	4 or 5*	A	*Higher test levels may be needed based on the installation location
Voltage Dips and Interruptions	61000-4-11 and/or 61000-4-29	4	A & B	Test System for capable of handling the voltage and current
Harmonic Immunity	61000-4-13	Class 3	A	Test System for capable of handling the voltage and current
Common Mode Disturbances	61000-4-16	4	A	CDN for voltage/current
Ripple	61000-4-17	4	A	Test System for capable of handling the voltage and current
Damped Oscillatory Waves	61000-4-18	4	B	CDN for voltage/current

11.3.3. DC Specific Immunity Gaps

Existing standards do not fully address DC bus transient disturbances, Converter interaction effects, Load-induced voltage fluctuations due to the conductive path self-inductance. It may be

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

of interest to consider Step-load transient testing, Bus disturbance injection testing, and Ripple immunity verification. These tests are critical for ensuring stable operation in dynamic load environments typical of data centers. The major cause of system instability will be the DC busway or conductors in the premises wiring DC distribution system itself.

11.4. Grounding and Bonding Strategies

Grounding is a critical determinant of EMC performance in LVDC systems.

11.4.1. Grounding Topologies

Grounding Systems	Application	Advantages	Challenges
Solid Grounded	Direct connection between DC return and earth	Stable voltage reference Effective noise suppression	Grounding loops Circulating currents Increased touch voltage to ground hazard for staff.
Floating	No intentional connection to earth	Eliminates ground loop currents	Increased common-mode noise
High-Impedance Grounded	Grounded reference through impedance	Controlled fault current Reduced noise compared to floating	Increased design complexity

11.4.2. Recommended Grounding Approach

While grounding the AC side is a common and workable approach for EMC, it must be done carefully. Floating the DC bus can reduce direct noise coupling and ground-loop currents, while grounding the AC output provides a stable reference for loads and helps protection devices function properly. However, the DC side is rarely truly “isolated” in practice, since EMI filters use

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

small capacitors to reference it to ground for noise control. The design must balance EMC performance with safety requirements, leakage current limits, and fault detection—especially in transformerless systems, where interactions are more complex. In short, it's a valid strategy, but it requires a coordinated grounding and filtering design rather than simply leaving the DC side completely unreferenced.

Sidecars separate the main power conversion path from the grounding and EMI control network by placing grounding-related components—like Y-capacitors (refer to 6.2.3.2 Y-Capacitor Limitations), common-mode chokes, and chassis bonds—in a parallel "sidecar" path connected to the enclosure or earth. Instead of directly bonding all nodes, the system keeps primary circuits (such as the DC bus) functionally floating or isolated while referencing them to earth through controlled, high-impedance paths. This approach helps manage common-mode noise, reduce EMI emissions, and avoid ground loops, while still providing a defined path for high-frequency currents and meeting safety requirements. It effectively balances EMC performance and electrical isolation by making grounding intentional and frequency-dependent rather than a simple hard connection.

11.4.3. Bonding and Return Path Design

Effective EMC requires that you have low-impedance bonding between racks and equipment, controlled return paths and minimizing loops. Multi-point bonding of the grounded conductor (Neutral conductor) is prohibited in the US National Electrical Code (NEC) NFPA 70-2026.

DC distribution impedance should be characterized and controlled (including cable inductance and busbar geometry) to prevent resonance, excessive ripple propagation, and instability under fast load transients.

11.5. Cable Shielding and Filtering

11.5.1. Cable Shielding

Unshielded cables are sufficient for AC input and DC input/distribution. It is preferable to use unshielded cables for I/O and control signals when possible but if needed for EMC/EMI performance shielded cables can be used but need to be grounded on both ends for effectiveness. When shielding cables you should avoid pigtail grounding.

11.5.2. Filtering

Both the AC input and DC output or input should be considered. On the input side you need to ensure they effectively limit the emissions on the power lines. On the DC output you want to

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

ensure that the power being distributed is not going to affect the equipment it is powering. Filters can be either common-mode or differential-mode and should consider damping networks to prevent resonance.

PSUs and power conversion stages used to supply loads inside the Racks need to be filtered at input side (towards the DC mains, 800V) and output side (towards the ELV, 50V).

11.6. System-Level EMC Considerations

11.6.1. Loading Considerations

Testing should reflect real world conditions the equipment could expect to see in use. It should be loaded for continuous operation with a load from 90-100%. The worst case EMI scenarios should be examined and applied based on the expected test conditions (this may mean that levels haven't been defined for certain tests, as well as expected performance criteria).

From an Emissions standpoint is it preferable to conduct testing at Low Load (5-10%, Medium Load - 50%, and High Load 90%)? Immunity testing should be conducted at 90% load.

In multi-converter LVDC systems (e.g., SST feeding multiple racks), system-level interaction testing should be performed to evaluate stability, oscillations, and EMI coupling between converters under dynamic load conditions.

11.6.2. Global Compliance

EMC requirements vary by region so for products that are deployed globally it is important to examine all the intended markets and then test to worst case limits and levels (where possible) to meet the necessary requirements. Common regional requirements:

- Europe (IEC / CISPR based standards)
- US (FCC)
- Asia-Pacific (CISPR based standards)

11.7. Gaps and Future Work

Currently there are areas that need to be clarified when it comes to designing and testing of LVDC Power Converters. Defining limits for Conducted Emissions, validating test levels and applicability of Immunity related DC power tests. Researching ripple/harmonic specifications for LVDC, and there is limited guidance on optimal grounding schemes. The major UPS

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

manufacturers all understand and have been dealing with DC power bus and battery issues for decades.

DC power system studies to quantify interactions between the multiple pieces of power electronic equipment, at the whole data center level. This is regularly done in AC power systems, it needs to be extended to DC.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

12. Codes and Standards

As direct current (DC) system architectures and microgrids gain momentum, especially in AI data centers where efficiency and reliability are critical for high-density computing, the role of global codes and standards (e.g., UL, IEC, IEEE, etc.) becomes increasingly important.

These codes and standards, along with the testing and certification protocols and demonstration of compliance that they empower, help enable installers, integrators, and operators to design and deploy safe, efficient, and interoperable DC systems.

However, the rapid growth of DC technologies is outpacing the standardization landscape, creating notable challenges.

This section of the white paper outlines the current landscape as well as gaps in the installation codes and standards for DC products and systems, explains why they exist, and outlines where industry cooperation is urgently needed.

This white paper continues to serve as an initial – and surely incomplete – accounting of the codes and standards that will support the transition to a higher voltage DC future. By identifying these challenges early, stakeholders can participate and accelerate safe innovation and help shape the next generation of DC codes and standards.

12.1. Why DC Codes & Standards Matter

- **Safety**: Clear regulations ensure that protection systems, fault management, and grounding schemes are consistently implemented.
- **Interoperability**: Harmonized standards allow components from different vendors to work seamlessly together.
- **Scalability & Future Expansion**: standardization simplifies system replication and modular growth, allowing capacity, voltage levels, or new loads to be added with minimal redesign, reduced risk, and lower costs over time.
- **Efficiency & Sustainability**: Proper DC design can reduce conversion losses.
- **Regulatory Compliance**: Standards enable local code and regulatory compliance, which helps streamline installation and commissioning.

12.2. Key Challenges in DC Standardization

1) **Alternating Current (AC) has a 100-Year Head Star**t: AC systems have benefited from decades of evolution, research, and global harmonization. In

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

contrast, DC is still emerging at scale – particularly at power levels relevant to data centers.

2) **Demand for DC Is Growing Faster Than Standards Can Evolve**: Industries want DC solutions now to meet efficiency and sustainability goals. But many DC-related standards are incomplete, outdated, or not globally aligned.
3) **Slow Standard Development Cycles**: Creating or updating standards typically takes several years, involving expert committees, testing cycles, and multinational coordination. This timeline does not match the pace of DC innovation and requires urgent action from Standards Development Organizations (SDOs) to accelerate the pace of requirements.
4) **Alignment Between International Standard Development Organizations (SDOs)**: Harmonization efforts between major bodies such as UL, IEC and IEEE are ongoing but slow. Misalignment leads to duplicated testing, engineering workarounds, and higher costs.
5) **Country-Specific Requirements**: Installation codes and regulations differ widely across regions, making global deployment complex. An installation compliant in Europe may require significant redesign to meet US or Chinese requirements.
6) **Cultural & Technical Inertia**: The industry is largely built around AC. Many professionals are reluctant to leave this comfort zone, creating skill gaps and hesitancy in adopting DC-based architectures. Training and education on DC systems will take time to implement and ensure workers are prepared.
7) **Utility Interface**: While the white paper primarily focuses on chip-to-grid power architecture, the point of common coupling (PCC) – where the data center meets the utility – presents significant integration hurdles. Data centers are not "static" loads; they are dynamic, high-density environments that strain traditional grid management in three specific ways: load swings, inrush and surge management, and grid balancing and curtailment. Without harmonized standards for utility-specific study requirements and commissioning criteria, operators face longer approval queues and higher "soft costs" before a single server is even racked.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

12.3. System-Level Standards

12.3.1. Existing system-level knowledge and applications

Multiple industries are already piloting, deploying, or operating DC microgrids at scale. These sectors provide valuable real-world experience that the data center ecosystem can leverage to accelerate adoption, reduce risk, and lower cost.

Industrial Applications: As outlined in previous sections of the white paper, factories, robotics, process automation, and traction increasingly use DC to improve efficiency, stabilize renewable integration, and simplify power conversion. These systems can offer insights into coordinated control, protection strategies, and energy flow management.

Renewables: Renewable energy sources like solar photovoltaics (PV) and fuel cells natively generate DC, and while wind turbines initially produce AC, their output is often rectified to DC for efficient control and storage. In modern applications, DC-coupled systems are emerging as the industry standard for co-located solar and battery storage.

Marine & Offshore: Ships and offshore platforms rely heavily on DC grids to optimize propulsion efficiency and manage distributed energy sources. Marine applications demonstrate mature approaches to system-level coordination and fault handling in compact, mission-critical environments.

EV Charging Infrastructure: In some regions, high-power EV fast chargers operate natively on DC, and large charging hubs use DC distribution for improved efficiency and simplified conversion. Their experience with high-current protection, conversion technologies, and the knowledge and learnings from AC/DC interfacing should be evaluated for transferability to data centers.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

Although their end applications differ, the fundamental structure of DC grids across these industries is remarkably similar:

- Distributed converters feeding a common DC bus
- Coordinated control of power flow and voltage levels
- High-power protection requirements
- Interfaces to AC utility grids or backup systems

This similarity creates a major opportunity. Data centers can learn, copy, and align with proven architectures rather than reinventing DC system concepts from scratch.

12.3.2. Existing system-level standards

A diverse coalition of Standards Development Organizations (SDOs) and industry-led alliances is currently engaged to codify the future of DC power.

Legacy safety leaders – including UL Solutions, NFPA, ANSI, IEC, and IEEE – are working in tandem with specialized trade bodies like the Open Direct Current Alliance (ODCA) and Current/OS to bridge the gap between traditional electrical safety and the unique demands of high-density DC architectures.

More than 20 system-level standards were identified for DC architecture through this white paper effort. Technical requirements are being reviewed for adaptation to data center applications, and several emerging technical reference documents are available to help planners, installers, and end users design and deploy DC systems.

Highlighted standards are below, and a full list of all identified system-level requirements applicable for DC power distribution architecture can be found in the master reference file located at this link.

- IEC 60364 Series (specifically IEC 60364-8-82), IEEE 2030.7 and 2030.8, IEEE 1547, UL 3001, UL 3001D, and Current/OS and ODCA system specifications – documents which define requirements of the DC system and grid interconnection.
- IEC 60038 (TR63282), IEC 60664-1, IEC 60364-4-44, and UL 840 – covering voltage and insulation coordination for equipment within low-voltage supply systems.
- UL 2941 – provides a comprehensive framework for testing and certifying the cybersecurity of grid-interactive equipment (inverters, BESS, and EV chargers).
- NFPA 75 – the standard for the Fire Protection of IT Equipment, which outlines requirements for construction, equipment, and fire suppression in data center environments.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

- ANSI/BICSI 002, ANSI/TIA-942C, IEC 22237, and Uptime Institute Tier Standard Topology – documents which define various classes for power redundancy and reliability applicable to electrical and associated cooling systems.
- IEEE has standards that apply to both AC and DC (IEEE C37.20), or some specifically for DC (IEEE 947); however, the industry expectation is that DC components must demonstrate performance parity with their AC counterparts in critical safety and resilience metrics, such as impulse withstand and short-circuit ratings. IEEE P3710.1 also has a document in development that will be a DC power system recommended practice for data centers from 300Vdc to 1500Vdc.

12.3.3. System-level gaps and standards in development

While the individual components for DC power distribution are maturing rapidly, the industry is navigating a governance gap regarding how these products function as a unified, resilient ecosystem. Rapid mobilization is required to address critical system-level blind spots that, if left unmanaged, could delay the scaling of AI-ready power distribution infrastructure. This effort focuses on moving beyond component safety to establish a holistic framework for architecture, control logic, and personnel protection. By codifying these emerging practices, the industry aims to replace custom, one-off engineering with repeatable, certified system designs.

12.3.3.1. **System Design Principles**: There is no unified guidance on how to architect a complete DC system and ongoing strategies are in development optimal distribution topologies, fault management strategies, grounding concepts, and coordinated protection schemes.

12.3.3.2. **System Control & Coordination**: Unlike AC systems, where frequency and voltage behavior are well defined, DC systems rely on control algorithms and coordinated converter behavior. Standards describing how devices should interact, share loads, and react during faults are still emerging. However, existing large-scale, commercial installations show the functionality and benefits of Low Voltage DC (LVDC) installations.

12.3.3.3. **Maintenance & Operational Procedures**: Operators lack standardized frameworks for inspection intervals, diagnostic procedures, and safe handling of energized DC equipment, especially considering unique DC hazards like sustained arcs.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

12.3.3.4. **AC/DC Interface Management**: The transition between AC and DC is critical but not yet fully regulated. Specifications for interoperability, protection coordination, and energy flow control at this interface vary widely across regions and applications.

12.3.3.5. **DC Arc Flash & Worker Protections:** Implementing DC power systems with increased voltage levels introduces critical gaps in NFPA 70E safety protocols, primarily because current standards and empirical models (like IEEE 1584) were developed for application to AC systems and more work is needed to ensure DC system risks are covered. While decades of rigorous research have enabled widely adopted AC arc-flash equations, no equivalent, comprehensive, and validated methodology exists for DC systems.

12.4. Product-Level Standards: Existing & In-Development

Product-level standards are crucial in providing the foundational safety, interoperability and performance requirements needed for effective design and deployment of components, products and sub-systems across DC system architectures.These need to cover nearly every component of the full system architecture across a range of product categories from chip-to-grid:

- Source
- Distribution Equipment
- Power Supplies/Converters
- Switching/Control Devices
- Overcurrent/Protection Devices
- Grounding/Bonding Equipment
- Connectors
- Cables/Cords
- Power Utilization
- Reference Standards

More than 70 product-level standards were identified for 800V DC architecture through this white paper effort. While a majority of these already address 800V DC requirements, a subset of gaps require standards to be updated, and in some cases, new standards need to be created.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

A full list of all identified product-level standards applicable for 800V DC power distribution architecture can be found in the master reference file located at this link.

12.4.1. Existing product-level standards

The majority of UL, IEC, and IEEE Standards across product categories for data center power have requirements that cover 800VDC, but the majority of products for these categories have not yet had products vetted or tested for 800V/1500V DC for data centers.

Some of these organizations are conducting gap analyses to determine required updates or revisions to address this issue. For example, all UL Standards with gaps not covering 800-1500VDC for data center use are in the process of being updated or created. To support the industry's need for speed and safety requirement adoptions, the UL Enterprise is working to expedite UL Standards updates with the following tools:

- **Certification Requirement Decisions (CRD)** - An official UL Solutions determination that establishes new or revised safety requirements for products, often facilitating the adoption of new technology before formal standards are updated. These decisions provide immediate compliance paths, often addressing specific technical, performance, or safety issues.
- **Outlines of Investigation (OOI)** - A non-consensus document, developed and owned by UL Solutions, containing specific requirements to certify products, components, or systems for which no published, formal standards currently exist. These documents are used to evaluate, test, and apply safety certification marks to emerging technologies or specialized products.
- **Category Control Numbers (CCN)** - A four-letter code assigned by UL Solutions to classify products and components based on their specific safety standards and intended use. These codes can be used to find guide information in the UL Product iQ database and allow for quick verification of safety compliance by regulators and manufacturers.
- **Guide Information** - UL Guide Information provides essential details on product safety standards, usage limitations, and installation requirements for UL-certified products, primarily accessed through the UL Product iQ database. It defines the scope of certification, applicable standards, and the specific Mark required for compliance. These guides cover categories like electrical safety, fire resistance, and building materials.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

12.4.2. Product-level standards in development

Many Standards Development Organizations (SDOs), including NFPA, IEEE, IEC, and UL Solutions are in a state of high mobilization to update their frameworks. The rapid shift toward DC systems has turned what were once slow review cycles into urgent, collaborative sprints.

As of early 2026, several key standards have either been updated or are in the final stages of revision to address high power DC systems. Again, an example of this exists where UL Solutions has engaged industry via meetings and working groups to address code updates, standards gap analyses, and exploration of emerging needs.

In addition, initial work has been done to evaluate and bring together the IEC standards portfolio, and more closely align with the power architecture approach being mapped using the portfolio of UL Standards. A LVDC Data Center Summit for the industry hosted by Current/OS took place in Aachen, Germany focused on addressing emerging needs in IEC standardization in late May 2026. Summit attendees included IEC leadership, chairpeople for key IEC technical committees, UL Solutions, and technical and standards leaders across multiple industries. A copy of the presentation can be found on the Current/OS website here.

Some gaps were identified where there was not an applicable IEC standard relative to the UL Standards for specific product categories such as Battery Energy Storage Systems (BESS, not directly comparable to UL 9540) and new product categories and standards being developed for data center cables, connectors, and other product category innovations. There were also different IEC interpretations for some of the electrical distribution product categories that will be further addressed.

A coordinated action plan is in development to identify a short-list of key IEC standards for 800V DC architectures, engage with those appropriate IEC committees, identify gaps and actions needed for those standards, and work towards expediting updates with some newer procedural tools within IEC. More information on this effort will be publicized in future white papers and made public at key data center industry conferences.

A brief selection of important standards in development are included here, and a more exhaustive list can be found in the master reference file located at this link.

Please note that since the publication of DCF Power Distribution LVDC white paper 1.0, several key areas have been identified to highlight updated information.

- IEEE SA P3105 – Recommended Practice for Design and Integration of Solid State Transformers in Electric Grid

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

- UL/IEC/CSA 62368-1 – Technical Committee (TC) 108 is exploring how to refine safeguards for high-power DC delivery. An adhoc group created within the TC108/HBSDT working group is drafting a requirements proposal to expand IEC 62368 scope to 1000VAC/1500VDC.

- UL 2877 for MV SST (also known as Power Units, Power Supplies)* – Medium Voltage Power Supplies are critical electronics based switch mode power supplies that are intended to step down medium voltage AC to DC at the facility/system infrastructure level. This future-state application will allow for fewer AC to DC power conversion stages throughout data center power architectures. They are sometimes referred to as medium voltage solid state transformers, with the category referred to as MV SST, but they are not actually transformers as the application will interface with the utility interconnection and are actually large-scale switch-mode power supplies. There is a lot of innovation happening in this space and this product category is addressed under UL 2877 OOI. Based on these innovations and applications for data centers, UL 2877 OOI is currently being updated to include bi-directional and other new application requirements. For integration with batteries, consider UL 1741 and UL 9540.

- Sidecars (also known as Power Racks) – multiple design iterations exist for sidecars/power racks and they continue to evolve based on power architecture system designs. Those iterations can include AC to AC, AC to DC, DC to DC, and different voltage levels. As this product category continues to evolve rapidly, contact a Nationally Recognized Testing Laboratory (NRTL, an OSHA-accredited designation) for appropriate standard application, requirements and evaluation of the product design.

- UL 2755 and UL 2140 (OOIs) – Modular data centers are a sub-system approach falling in-between an individual product and a full system architecture. It differs from field installation system assembly by incorporating integrations of multiple equipment types into a prefabricated assembly. This approach provides multiple benefits to system designers, data center owners and manufacturers. These modular architectures can be containerized modules, which are addressed under the recently-updated UL 2755 Version 3.1 OOI, or open power skids, which are addressed under UL 2140 OOI.

  Different types of module assemblies include:

  - Power modules

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

- IT modules
- Cooling modules
- Hybrid modules
- Engine Generators (GenSets), open or enclosed types, which fall under UL 2200/UL 6200

There is significant activity in the modular data center space with 2029 NEC code submissions and an IEEE Modular Data Center Design Guide in development (P3710).

- Ground fault / insulation monitoring
  - For solidly grounded systems, UL is developing DC GFPE and GFCI requirements under the Outline of Investigation process, while IEC DC-RCD requirements do not presently extend to 800V and are in discussion for development.
  - For impedance grounded (high resistance grounded) systems, UL is developing Neutral Grounding Resistor (NGR) monitoring requirements under the Outline of Investigation process, while IEEE 32 requirements exist for the resistors themselves and CSA C22.1 (Rule 10-302) requirements for monitoring also exist. For floating (ungrounded) systems, UL is developing monitoring requirements under the Outline of Investigation process, while IEC has already developed 61557-1 general requirements, 61557-8 for Insulation Monitoring Devices (IMDs), 61557-9 for Insulation Fault Location System (IFLS), and is developing 61557-19 for protective earth monitoring.

- NFPA 1 / NFPA 855 / UL 1973 / UL 9540 / UL 9540A for ESS – The applicability of existing energy storage system (ESS) codes and standards is being challenged by the rapid evolution of datacenter architectures. Historically, frameworks such as NFPA 855 and UL 9540A have been developed around centralized, large-scale installations with clearly defined system boundaries, spatial separations, and hazard mitigation approaches. Emerging designs with ESS closer to the load underscores the need for coordinated standards development to address distributed, multifunctional storage architectures and ensure consistent safety requirements that protect personnel and increasingly valuable datacenter infrastructure.
- Cybersecurity – a critical consideration for low-voltage DC electrical distribution systems, particularly as these systems incorporate digital controls, monitoring capabilities, and connectivity to external and local communication networks. These connections, whether through Ethernet, wireless interfaces, building management systems (BMS), industrial

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

control systems (ICS), or cloud-based platforms, introduce potential pathways for unauthorized access, intrusion, or malicious attacks.

Unlike traditional isolated electrical systems, connected DC distribution systems may include intelligent electronic devices (IEDs), programmable logic controllers (PLCs), power converters, battery management systems, and monitoring platforms. If compromised, these components can be intentionally or unintentionally manipulated to alter operational parameters such as voltage levels, current limits, or load prioritization. Such actions can directly impact system reliability and, more importantly, safety. For example, malicious control could lead to overheating conductors, improper operation of protective devices, battery overcharging or deep discharge, or unexpected shutdown of critical loads.

The consequences of a cybersecurity breach extend beyond equipment damage to potential risks to human safety, service continuity, and mission-critical operations. Addressing cybersecurity through secure product design, adherence to recognized standards such as ANSI/ISA/IEC 62443-4-1, and comprehensive system-level risk assessments per ANSI/ISA/IEC 62443-3-2, are necessary to mitigate these risks. Additional standards and references to consider include UL 2900, UL 2941, IEEE 1547.3, and NEMA CY70001-2023.

12.5. Installation Codes: Existing & Gaps

DC power distribution in data centers is governed by a combination of regional- and/or country-level electrical safety codes, fire protection standards, and industry best practices. Applicable commissioning standards should be consulted and need to be met for each region/installation location.

For electrical installations and safety-related aspects, there are two main Standards Development Organizations: International Electrotechnical Commission (IEC) and Underwriters Laboratories (UL). Whereas UL standards target primarily the United States and Canada (North American markets), IEC standards target global markets, especially Europe and Asia. In addition to the difference in the geographical scope, IEC standards focus on minimum safety and performance requirements for electrical products, while UL Standards offer stricter, more detailed descriptions and rigorously focus on in-depth product testing, factory inspections, and follow-up audits.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

This white paper provides a summary of global IEC-based installation requirements below, while the US is covered in more specific detail. Deeper work is required to provide a comprehensive global analysis on the current and emerging code and regulatory landscape. This will be covered in a future revision of this document.

12.5.1. Global Installation Requirements

As noted in the Voltage section above, in the context of power distribution, low voltage is defined by international product standards and regulations along with installation codes as operation at or below 1000 VAC or 1500 VDC.

Most countries outside the United States (US) use IEC 60364 as the base for installation rules with local translation and modifications.

IEC Technical Committee 64 (IEC TC 64) is responsible for the IEC 60364 series, the international "gold standard" for low-voltage electrical installations. These rules provide a comprehensive framework for the design, erection, and verification of electrical systems to ensure the safety of persons, livestock, and property.

This applies to installations with nominal voltages up to 1000VAC or 1500VDC. The installation rules are organized into several key parts:

- **Part 1: Fundamental Principles** – Defines the scope, basic safety requirements, and general characteristics of an installation.
- **Part 4: Protection for Safety** – This is the "safety engine" of the rules, covering:
    - **Electric Shock**: Requirements for automatic disconnection of supply, insulation, and the use of Residual Current Devices (RCDs).
    - **Thermal Effects**: Rules to prevent fires caused by electrical equipment.
    - **Overcurrent**: Protection against short circuits and overloads.
    - **Voltage Disturbances**: Protection against surges and electromagnetic interference.
- **Part 5: Selection and Erection of Equipment** – Guidance on choosing and installing wiring systems, switchgear, and earthing arrangements.
- **Part 6: Verification** – Specifies mandatory initial inspections and periodic testing to confirm the installation remains safe over time.
- **Part 7: Special Installations** – Tailored rules for high-risk locations like bathrooms, swimming pools, medical facilities, and solar photovoltaic (PV) systems.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

- **Part 8: Functional Aspects** – Modern additions focusing on energy efficiency and "prosumer" installations (where buildings both consume and produce energy).

12.5.2. United States Installation Requirements

In the US, DC systems must operate within a hierarchy of federal safety standards, construction codes promulgated by various levels of government including state and municipal agencies responsible for building safety and industry-specific fire protection requirements that are developed and managed by **NFPA 70® (National Electrical Code®)**:

- Applies to AC and DC premises wiring systems connected to the load side of the utility service point. Current requirements are not voltage or system limited. Most current DC requirements are general in nature and data center specific infrastructure requirements do not currently exist. The next revision will be completed in the third quarter of 2028. There are task groups focused on Bi-Directional Power, Data Centers, and DC Systems, which are reviewing and identifying gaps for DC systems in general and AC & DC infrastructure specific to data centers.
- The National Electrical Code® developed by NFPA® is revised on a three-year schedule and is incorporated into statute or administrative rule by state and/or local governmental agencies in all 50 US states and several international jurisdictions. None of these agencies automatically update the NEC® when it is revised and the current landscape of editions in effect ranges from the 2008 to the 2023 editions. Like the code development process, there is a process by which these agencies to promulgate new construction and electrical safety codes and these processes often involve multiple steps of review and public comment which naturally delays implementation of the most recent electrical safety requirements. Learn where the NEC is enforced.
- National Fire Protection Association (NFPA) standards, US Occupational Health & Safety Administration (OSHA), and the US Department of Energy (DOE). Current requirements focus primarily on AC systems, so specific sections of code and related standards are actively under evaluation to address the unique hazards of DC systems.

- **Existing Codes**
  - **NFPA 70® (National Electrical Code®-2026):** Applies to AC and DC premises wiring systems connected to the load side of the utility service point. Current requirements are not voltage or system limited unless otherwise specified. Most

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

current DC requirements are general in nature and data center specific infrastructure requirements do not currently exist. The next revision will be completed in the third quarter of 2028 and there are task groups reviewing and identifying gaps for DC systems in general and DC infrastructure specific to data centers. www.nfpa.org/70

- **NFPA 70B (Standard for Electrical Equipment Maintenance-2026)**: Requirements that detail the preventive maintenance of electrical, electronic, and communication systems. Covers AC and DC equipment, but more focus on AC infrastructure. Has synergy with NFPA 70 and NFPA 70E personnel safety requirements that contemplate electrical equipment being properly maintained. Voluntary compliance is primary use, not mandated like the NEC®. The next revision cycle is for the 2029 edition. www.nfpa.org/70B
- **NFPA 70E® (Standard for Electrical Safety in the Workplace-2024)**: Focuses on personnel safety, specifically shock and arc flash hazard protection. Applies to AC and DC systems rated over 50 volts. Is promulgated by US DOE and is used by the US Department of Labor, OSHA in general duty clause citations. More focus on AC systems. Currently near completion for the 2027 revision cycle. The next revision cycle for 2030 edition. www.nfpa.org/70E
- **NFPA 75 (Standard for the Fire Protection of Information Technology Equipment-2024)**: Standard contains requirements for the protection of information technology equipment and information technology equipment areas from damage by fire or its associated effects-smoke, corrosion, heat, and water. The next revision cycle is for the 2028 edition. www.nfpa.org/75
- **NFPA 855 (Standard for the Installation of Stationary Energy Storage Systems-2026)**: Standard provides the minimum requirements for mitigating hazards associated with energy storage systems. Standard applies to utility and customer owned ESSs. The next revision cycle is for the 2029 edition. www.nfpa.org/855
- **NFPA 37 (Standard for the Installation and Use of Stationary Combustion Engines and Gas Turbines-2024):** Standard establishes criteria for minimizing the fire hazards related to the installation and operation of stationary combustion engines and stationary gas turbines that are fueled by liquid or gaseous fuels and are used as the prime movers for emergency generators, fire pumps, and

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

stand-by and peak power systems. The next revision cycle is for the 2027 edition. www.nfpa.org/37

- **NFPA 1(Fire Code-2024)**: Code advances fire and life safety for the public and first responders as well as property protection by providing a comprehensive, integrated approach to fire code regulation and hazard management. It addresses all facets of fire protection with extracts and references to over 130 NFPA codes and standards. Has many system and occupancy based chapters. There is a chapter on telecommunications facilities and information technology equipment but no chapter specific to data centers. It is promulgated in a number of US states and municipalities. The next revision cycle is for the 2027 edition. www.nfpa.org/1
- **International Fire Code (IFC)** – This is a comprehensive model code published by the International Code Council (ICC) to establish minimum regulations for fire prevention, life safety, and property protection in new and existing buildings. It regulates hazardous materials, fire protection systems, and emergency responder access to minimize risks from fire, explosions, or hazardous material discharges. There is no specific chapter for data centers. The IFC is promulgated in a number of US states and municipalities. The current revision cycle is for the 2027 edition. https://codes.iccsafe.org/content/IFC2024V2.0

12.5.3. Emerging trends and recommendations

When all these elements come together, how will local authorities handle commissioning and approvals; DC vs. AC requires training and education.

- **US & Canada** – State & Federal Jurisdiction: Large-scale data center power interconnections (often exceeding 20MW) are increasingly coming under the scrutiny of the Federal Energy Regulatory Commission (FERC) and state-level bodies (e.g., Texas SB-6) regarding grid reliability and cost allocation. **Energy Performance Standards**: Newer regulations, such as **California's A.B. 222**, may soon require data center operators to report energy consumption and adhere to efficiency performance standards.

12.6. Interim Guidance

DC technology is advancing faster than system- and product-level standards can adapt. While many product standards are available, and some system standards are emerging, guidelines for

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

designing, controlling, maintaining, and integrating DC systems remain under development. Where gaps in standards or regulation exist, OEMs can still prepare and mitigate risk by taking thoughtful steps to help assure the safety and performance of new technologies and end-use applications.

- **Evaluate Existing Standards** – Review standards, OOIs, or CRDs where a close alignment to the technology and risks can be demonstrated. For example, standards written for high-power DC systems in photovoltaic (PV) and electric vehicle (EV) applications can be evaluated for learnings, risks, and testing considerations.
- **Hazard-Based Risk Assessment** – Engage a certifying body and use a hazard based risk assessment approach to conduct pre-certification validation, interim testing, or field evaluation.
- **Modeling and Simulation** – Hardware-in-the-Loop (HIL): PHIL/CHIL, and EMT
- **Laboratory Testing** – high power laboratory testing for isolation, protection strategies, power quality performance
- **Pilot Projects** – potential benefits include faster pilot deployment/time to data, lower costs, faster validation of reliability, controls, and communications before advancing to higher voltage classes.
    - Bidirectional power flow control, performance under grid faults.
    - DC bus stabilization.
    - Utility communications (and sometimes curtailability).
    - Storage at different locations:
        - GPU serverboard
        - Rack DC bus (100’s ms-sec) GPU load peaks/valleys, training power smoothing
        - Row/PDU/UPS - fault ride through
            - Online ESS/UPS - seconds-minutes
            - ESS/UPS - minutes/tens of minutes
        - Energy resilience/grid services - minutes to hours
            - DC-coupled storage/generation
            - AC-coupled storage/generation

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

13. The path forward - Key Takeaways for Stakeholders

The transition to high-density DC architectures in AI data centers is no longer a theoretical pursuit – it is a technical necessity. As this white paper has outlined, the efficiency, scalability, and reliability gains offered by DC systems are clear. However, the path to global adoption is currently constricted by a standardization landscape that was built for an AC world.

While AC systems benefited from a century of refinement, the DC ecosystem is being forced to mature in a fraction of that time. We are operating in a "standardization gap" where innovation is outstripping formal regulation. This has created a unique set of challenges and opportunities, from the lack of unified guidance on system-level topologies to the urgent need for a validated DC arc flash methodology.

- **The Foundation Exists:** Despite gaps, existing product-level standards and proven architectures from the marine, industrial, and EV sectors can serve as a launching pad for DC design and deployment.
- **Collaboration is Non-Negotiable**: The "Known Unknowns" of DC technology, such as complex utility interfaces and high-power protection, cannot be solved in isolation. Success requires active participation and alignment between SDOs like UL, IEC, IEEE and NFPA.
- **Proactive Risk Mitigation**: Until formal installation codes like the 2029 NEC® are finalized, organizations must lean on interim guidance, utilizing Outlines of Investigation (OOI), conducting hazard-based risk assessments, and leveraging Hardware-in-the-Loop (HIL) simulations.
- **Voltage and Power Quality Alignment:** Key focus for further work to align the industry and further the discussion specifically on scaling. Sub workstream focus is needed for the next version of the paper.
- **DC Arcflash Rules:** Need work to be focused on and aligned as an industry as this analysis as it is the current bottleneck for the supply chain development and installation for the direct current infrastructure.
- **Power Distribution - Rack and Power Alignment:** A sub-workstream for the next version of the paper to fully understand the handoff to support collaboration and acceleration of the industry.

By closing the gap between the pace of innovation and the pace of regulation, we can ensure that the next generation data center power infrastructure is not only faster and more efficient but

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

fundamentally safer. The frameworks under development today will serve as the bedrock for a decentralized, DC-powered future.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

14. Conclusion

The shift towards AI-driven workloads is fundamentally reshaping the data center power landscape, demanding greater efficiency and scalability than traditional AC architectures can easily accommodate. This white paper, a collaborative effort within The Open Compute Project, establishes a foundational technical framework for understanding and evaluating direct current (DC) power distribution systems in data center facilities. We have systematically examined the architectural concepts, key design trade-offs – including voltage selection, polarity, system earthing, and protection schemes – and the critical interplay with main system components. The document purposely presents representative DC approaches observed across the industry, rather than mandating a singular solution, to foster an environment of open discussion and consistent interpretation among all stakeholders, including operators, designers, manufacturers, utilities, and AHJs. By aligning with the OCP tenets of openness, efficiency, impact, scale, and sustainability, this work aims to facilitate informed decision-making as DC architectures continue to evolve. While this paper provides a robust initial framework, it concurrently identifies specific areas for further collaborative development to ensure the path forward results in safe, efficient, and widely-adopted next-generation power solutions. The ongoing convergence of energy and IT infrastructure will continue to drive a multi-modal power environment where both AC and DC models must coexist and evolve.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

15. References

15.1. Papers

- [Ca] P. Cairoli, R. Rodrigues and H. Zheng, "Fault current limiting power converters for protection of DC microgrids," SoutheastCon 2017, Concord, NC, USA, 2017, pp. 1-7, doi: 10.1109/SECON.2017.7925392.
- [Mar] A. Marroquin, C. Carne, W. Brown and T. Landry, "Transient DC Arc-Flash Incident Energy Calculations for DC Distribution Systems," 2025 IEEE IAS Electrical Safety Workshop (ESW), Jacksonville, FL, USA, 2025, pp. 1-12, doi: 10.1109/ESW58401.2025.11159088.
- [Rav] S. Ravyts, G. V. d. Broeck, L. Hallemans, M. D. Vecchia and J. Driesen, "Fuse-Based Short-Circuit Protection of Converter Controlled Low-Voltage DC Grids," in IEEE Transactions on Power Electronics, vol. 35, no. 11, pp. 11694-11706, Nov. 2020, doi: 10.1109/TPEL.2020.2988087.
- [Sh] J. J. Shea, T. Landry and M. Liptak, "Short-Circuit Faults in DC Microgrids," 2024 IEEE Sixth International Conference on DC Microgrids (ICDCM), Columbia, SC, USA, 2024, pp. 1-6, doi: 10.1109/ICDCM60322.2024.10665093.
- [ODCA], Johann Austermann, Hartwig Stammberger [Ed]: “LVDC Microgrid system description”, 13. Apr. 2026, retrieved 05 June 2026: https://odca.zvei.org/resources/publications/update-of-odca-system-description-for-dc-microgrids-april-2026.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

15.2. Web Resources

- [1] Choukse, Esha, et al. "Power stabilization for AI training datacenters." arXiv preprint arXiv:2508.14318 (2025).
- [2] S. Bartlett, et al. (Oct 2, 2024). Data Center Facility Sustainability Metrics. OCP White Paper. [Online]. Available: https://www.opencompute.org/documents/dcf-sustainability-metrics-final-r3-docx-pdf.
- [3] A. Desimone, et al. (July 2025). Embodied Carbon Disclosure Base Specification v1. OCP Base Specification. [Online]. Available: https://www.opencompute.org/documents/carbon-disclosure-ocp-base-specification-20250730-pdf.
- [4] M. Dansie, et al. (Dec 2024). Requirements for Energy Storage Systems used in Data Centers v1. OCP White Paper. [Online]. Available: https://docs.google.com/document/d/1lQfdPtMe-XVQzydZVKTO6mnoubHUkn1UjBAa70Cbx-Q/edit?tab=t.0.
- [5] G. Yang and M. Tu (Dec 2025). Protections for Data Centers Powered by Direct Current. [Online]. Available: https://assets.new.siemens.com/siemens/assets/api/uuid:cacbc29b-5d0d-4b12-9053-696eb0b43e74/Protections-for-Data-Centers-Powered-by-Direct-Current-Technical-Paper_original.pdf.
- [6] Helmut Pusch, Schaltbau: “2.5 years experience of DC Microgrid operation in industry”, 11. Dec. 2025, https://news.schaltbau.com/schaltbaus-next-factory-lessons-from-the-worlds-first-dc-microgrid-powered-plant/#.
- [7] Tobias Lüke, Iris Schaper: “Path to Sustainability in car manufacturing with Direct Current (German title: Mit Gleichstrom zu mehr Nachhaltigkeit in der Automobilproduktion)”, 22. Oct. 2024, https://kem.industrie.de/elektrotechnik/mit-gleichstrom-zu-mehr-nachhaltigkeit-in-der-automobilproduktion/
- [8] ABB press release: “World’s most powerful DC data center online, 30. May 2012, https://new.abb.com/news/detail/12816/worlds-most-powerful-dc-data-center-online.
- [9] Hartwig Stammberger, ODCA, OCP EMEA Summit 2026: “Next Gen AI data centers with industry-proven direct current power supply”, 29. Apr. 2026, available:

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

https://drive.google.com/file/d/1TQcs9Ij3KrkNjJIrkBHk6RKLx--hffms/view?usp=drivesdk on 2026 OCP EMEA Summit » Open Compute Project.

- [10] Incident Review – Considering Simultaneous Voltage-Sensitive Load Reductions (8 Jan 2025) https://www.nerc.com/globalassets/our-work/reports/event-reports/incident_review_large_load_loss.pdf
- [11] NOGRR282 – Large Computational Load Ride-Through Requirements (status: Pending at cut-off) https://www.ercot.com/mktrules/issues/NOGRR282
- [12] SPP HILL Fault Ride Through Requirements, Version 1.0; RR696 – Integrate & Operate High Impact Large Loads https://www.ercot.com/mktrules/issues/NOGRR282
- [13] SPP HILL Fault Ride Through Requirements, Version 1.0; RR696 – Integrate & Operate High Impact Large Loads https://www.spp.org/documents/74635/spp%20hill%20fault%20ride%20through%20requirements-v1.0-8-19-2025%20%28updated%20version%29.docx
  RR696 / HILL integration page: https://www.spp.org/markets-operations/high-impact-large-load-hill-integration/
- [14] Technical Requirements for Large Computational Loads Connecting to the Ontario Power System, Version 1.0 (new requirements proposed for incorporation after stakeholdering) https://ieso.ca/-/media/Files/IESO/Document-Library/engage/large-computational-loads/LargeComputationalLoads-20260501-Technical-Requirements.pdf
- [15] MPID345 Demand Facilities FRT Recommendation Paper (1 Apr 2026; submitted for approval) https://cms.eirgrid.ie/sites/default/files/publications/MPID345-Demand-Facilities-FRT-Recommendation-Paper-01042026.pdf
- [16] Diablo 400 Specification v0.7.0 (1 March 2026) https://www.opencompute.org/documents/ocp-specification-diablo-400-v0-7-0-final-pdf
- [17] M. Zahedi et al. (Quanta Technology) — Best Practices for Large Load Interconnections. CIGRE 2025 (arXiv:2601.12686).
- [18] NERC — Characterization and Risks of Emerging Large Loads (July 2025).
- [19] NERC — Reliability Guideline: Risk Mitigation for Emerging Large Loads (May 2026).

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

- [20] ERCOT — Nodal Operating Guide Revision Request (NOGRR) 282 (2026).
- [21] SPP — High Impact Large Load Ride-Through Requirements, v1.0 (Aug 2025), Table 1.
- [22] ATC — Transmission System Planning Criteria PLG-CR-0001-V25 (28 Aug 2025), §9.2 Table 9.2-1. https://cdn.misoenergy.org/ATC%20TO%20Planning%20Criteria108210.pdf
- [23] IESO — Technical Requirements for Large Computational Loads Connecting to the Ontario Power System, v1.0 (1 May 2026), Table 1 (+ frequency table).
- [24] AESO — Guide to AESO Connection Requirements for Transmission-Connected Data Centres (12 Jun 2026), §6.1, Tables 6-1 and 6-3. https://www.aeso.ca/assets/Uploads/grid/Guide-to-AESO-Connection-Requirements-for-Transmission-Connected-Data-Centres.pdf
- [25] EirGrid/SONI — MPID345 Large Demand Facility FRT: Information Paper (Nov 2025, §6.6 Fig 18); Demand Facilities FRT Recommendation Paper (1 Apr 2026, Grid Code mod. CC.7.4.3). https://cms.eirgrid.ie/sites/default/files/publications/MPID345-Demand-Facilities-FRT-Recommendation-Paper-01042026.pdf
- [26] Energinet — Technical Regulation 3.4.3, Rev 1 (eff. 1 Nov 2024).
- [27] RTE — Documentation Technique de Référence, Article 8.3.5 v1.0 (eff. 15 Feb 2024), §3.12.
- [28] Fingrid — Technical requirements for large demand facilities (12 Jun 2026), §9.1.10.
- [29] European Union — Network Code on Demand Connection, Regulation (EU) 2016/1388.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

16. License

OCP encourages participants to share their proposals, specifications and designs with the community. This is to promote openness and encourage continuous and open feedback. It is important to remember that by providing feedback for any such documents, whether in written or verbal form, that the contributor or the contributor's organization grants OCP and its members irrevocable right to use this feedback for any purpose without any further obligation.

16.1. Creative Commons

It is acknowledged that any such documentation and any ancillary materials that are provided to OCP in connection with this document, including without limitation any white papers, articles, photographs, studies, diagrams, contact information (together, “Materials”) are made available under the Creative Commons Attribution-ShareAlike 4.0 International License found here: https://creativecommons.org/licenses/by-sa/4.0/, or any later version, and without limiting the foregoing, OCP may make the Materials available under such terms.

As a contributor to this document, all members represent that they have the authority to grant the rights and licenses herein. They further represent and warrant that the Materials do not and will not violate the copyrights or misappropriate the trade secret rights of any third party, including without limitation rights in intellectual property. The contributor(s) also represent that, to the extent the Materials include materials protected by copyright or trade secret rights that are owned or created by any third-party, they have obtained permission for its use consistent with the foregoing. They will provide OCP evidence of such permission upon OCP’s request. This document and any "Materials" are published on the respective project's wiki page and are open to the public in accordance with OCP's Bylaws and IP Policy. This can be found at http://www.opencompute.org/participate/legal-documents/. If you have any questions, please contact OCP.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

17. About Open Compute Foundation

The Open Compute Project (OCP) brings at-scale innovations and hyperscaler best practices to all, spanning technology domains from the data center to the edge, and the technology stack from silicon, to systems, to site facilities and services. The international OCP Community is made up of organizations and people from hyperscale and tier-2 cloud data center operators, communications providers, colocation providers, diverse enterprises, and technology vendors. The OCP Foundation fosters collaboration within the Community to tackle market challenges in an array of OCP Projects that create open contributions such as specifications, designs, and more. The OCP Solution Provider Program then recognizes the application of those contributions in compliant products, solutions, and data center facilities and services with designations like OCP Inspired, OCP Accepted and OCP Ready. With the tenets of openness, impact, efficiency, scale and sustainability, the OCP engages and educates thousands of engineers every year through many events and webinars. Across many initiatives the OCP Foundation and Community are meeting the market today and shaping the future.
Learn more at: www.opencompute.org.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.