Investigative Report on Electrical Commissioning in HPC Data Centers

Joseph Prisco
IBM
prisco@us.ibm.com

Grant Stewart
Los Alamos National Laboratory
gstewart@lanl.gov

Herbert Huber
Leibniz Supercomputing Center (LRZ)
herbert.huber@lrz.de

Randy Rannow
Silverdraft Supercomputing
pi-boson@ieee.org

Jason Hick
Los Alamos National Laboratory
jhick@lanl.gov

Dave Martinez
Sandia National Laboratory
davmart@sandia.gov

Brandon Hong
Lawrence Livermore National Laboratory
hong13@llnl.gov

Aditya M. Deshpande
Microsoft Corporation
aditya.deshpande@microsoft.com

***Abstract* – The Energy Efficient High Performance Computing Working Group (EE HPC WG) has assembled a small diverse team to write a short investigative report on electrical commissioning. The purpose of the investigative report is to evaluate the need for electrical commissioning guidelines specific to High Performance Computing (HPC) data centers given their unique IT equipment load densities and power profiles. It is the consensus of the team that special electrical commissioning guidelines are needed and the EE HPC WG will author the initial guidelines. The scope of the guidelines will include the static and dynamic electrical aspects of commissioning practices that are specific to high performance computing and more importantly, cover the transient aspects of electrical commissioning. The fluctuating nature of many compute nodes can dramatically influence generation, transmission, and distribution of electrical power. HPC data center lessons learned and best practices will be examined and used to enhance the electrical commissioning guidelines. The primary audience for the guidelines is facility engineers and operators of HPC data centers. The guidelines will also be applicable to others that support HPC data centers, ranging from utilities and their electrical grid infrastructure to IT equipment manufacturers whose machines are being commissioned at the end of the process.**

Keywords – commissioning, electrical infrastructure, high performance computing

I. Introduction

The Energy Efficient High Performance Computing Working Group (EE HPC WG) created an Electrical System Commissioning Team to evaluate whether or not there is a need for electrical commissioning documents for HPC data centers. This short high level investigative report identifies the need for electrical commissioning documents that account for the unique elements of HPC data centers.

Commissioning is a "systematic process of verifying that systems within the built environment, beyond a building's skin, perform in accordance with design intent and the property owner's operational needs [1]." Commissioning is not only for new construction of facilities and data centers, but should also be done for any project, including renovations and retrofits to existing buildings where electrical infrastructure components may have been deployed decades earlier and may now be unable to meet the intended need. Commissioning should be done from inception to completion of a data center facility project. Electrical commissioning standards and handbooks from organizations such as InterNational Electrical Testing Association (NETA) and Building Commissioning Association (BCxA) cover the topic in much greater depth than in this short investigative report. There are many opportunities for the commissioning engineer to influence a project in the submittal, design, construction, acceptance, and occupancy phases.

HPC data center commissioning is an important set of actions that can minimize the risk of unplanned outages and downtime. According to the Building Commissioning Association (BCxA), "…the basic purpose of commissioning is to provide documented confirmation that building systems and assemblies function in compliance with criteria set forth in the Project Documents to satisfy the owner's operational needs [2]." The Owner's Project Requirements (OPR) is the main document that provides the cost, performance, and operational goals of the building whether it is general office space or a mission critical facility. When commissioning a mission critical facility such as a data center, reliability is the primary goal given the importance of uptime and the value of systems. Because many data centers cannot tolerate single points of

978-1-7281-6677-3/20/$31.00 ©2020 IEEE
DOI 10.1109/CLUSTER49012.2020.00074

failure without the risk of downtime or damage to systems, reliability is often accomplished via multi-level redundant solutions to sustain availability objectives.

Other goals of commissioning may include U.S. Green Building Council (USGBC) Leadership in Energy and Environmental Design (LEED) certification as well as immunity to climate change. LEED certification is a green building rating system that provides a framework for energy efficiency and cost savings. LEED BD+C specifically targets data centers, recognizing improvements to efficiency in these energy-intense buildings. Climate change is impacting the natural operating environment of utilities and data centers around the world. Simulating the impacts of individual events or a series of cascading events through commissioning allows weaknesses to be identified and corrected, ultimately hardening the data center's electrical infrastructure.

It is appropriate to commission individual pieces of equipment and entire systems. This is required to ensure satisfactory operation of the equipment as a stand-alone entity, as well as when assembled into a larger functional environment with complex controls. Although it is very time consuming and costly to fully test all equipment and systems in a data center to cover all possible failure scenarios, the commissioning process must ensure that when a data center is turned over to the owner, there is a high level of confidence that the desired level of reliability, redundancy, and availability can be achieved. It is difficult to put a price tag on deficiencies found during commissioning which are not quantifiable, but could save energy, lives, and avert unplanned outages as well as downtime [3].

There are several major categories for commissioning, including, but not limited to mechanical (HVAC), electrical, plumbing, life safety, and the building structure. Mechanical, electrical, and plumbing is often lumped together and abbreviated MEP. The most attention is typically given to the mechanical plants supporting data centers because of the required temperature, relative humidity, and dew point for the IT equipment. However, commissioning of the electrical system is usually the longest duration and results in many electric power on and off cycles [4]. During the electrical testing, it is difficult to commission other systems without stable electrical power.

II. Electrical Commissioning

Besides the facility and IT equipment, ancillary data center equipment in the mechanical plant (e.g., chillers, cooling tower fans), plumbing (e.g., pumps, valves), and life safety (dampers, alarms, fire suppression) all require electricity to operate. The electricity must be reliable and of acceptable voltage regulation and quality.

Due to the individual design and construction aspects, not all data centers will follow a specific electrical commissioning playbook. Many times, it is up to the commissioning agent in consultation with the facility owner, general contractor, architect, engineers, and consultants to perform commissioning as outlined in the OPR and Basis of Design (BoD) documents.

Electrical commissioning may consist of several basic tasks, all aimed at finding potential problems and identifying fixes to keep the design and construction on schedule and on budget. The submittal or procurement phase may insist in request for proposals (RFPs) and tenders that commissioning be included by a dedicated agent. During the design phase, electrical drawings and specifications will be reviewed to catch issues early where the fix can be easily made. Simulations may be used to analyze typical operation and fault scenarios where the interaction between various pieces of electrical equipment is warranted. Other studies such as selective coordination and fault current propagation may also be done.

Equipment start-up will usually occur toward the latter stages of construction. Electricity is needed to ensure the equipment can power on correctly based on the manufacturer user's guide and instructions. After the construction phase is the acceptance phase where the operation of equipment and compliance verification is confirmed by witness testing. Factory testing prior to shipment of equipment to the job site is typically done by the manufacturer, but on-site testing confirms that the equipment survived transportation and is installed correctly. Upon successful start-up, equipment operation is verified as an individual unit and interactively as part of a larger system. Normal operations can be checked with load banks to simulate IT equipment electrical input and heat output [5]. This includes generators, uninterruptible power supplies (UPSs), control systems, chillers, and computer room air conditioners (CRACs) as well as computer room air handlers (CRAHs). Load banks can be used to load circuits and equipment. While under load, infrared imaging is used to scan electrical connections and to identify hot spots that need attention (e.g., proper terminations, torqueing). Emergency conditions are simulated to mimic individual equipment failure redundancy, loss of local power (e.g., power panel, switchgear) and site power backup capability, and maintenance modes. Staff training is also done during the acceptance phase. On-site personnel need to be trained to understand how to operate and test the system.

The last electrical commissioning task is the occupancy phase where the baseline state of the installation is documented and systems manuals, warranty documentation, and test reports are turned over to the owner.

III. HPC Data Centers

High performance computing data centers are like any other enterprise or cloud data center with servers, storage, and networking hardware installed. However, HPC resources and services support customers who solve large, complex engineering and scientific problems, visualize and interpret results, and manage and retain large volumes of data. To accomplish these tasks, HPC data centers generally have large clusters of compute and storage nodes that act as large block loads. This arrangement is different than enterprise or cloud data centers where each computing node can start or stop individually. Because the compute nodes are aggregated for parallel processing and act in unison, load fluctuations are more extreme in HPC data centers compared to enterprise data centers. HPC capability in a data center requires electrical power and cooling infrastructure that supports very fast electrical transients (di/dt) and clusters of many high-density kW racks.

The basic electrical commissioning of an HPC data center should be like enterprise or cloud data centers, but include additional design verification and testing that account for the sustained performance and power consumption of flagship jobs. Electrical commissioning procedures are needed for new HPC systems and expansion of HPC systems due to their high capacity and density as well as their impact on the stability of a building's electrical distribution system. Commissioning should involve validating the HPC systems electrical distribution capabilities under high user application load to determine acceptable capacity for application workload, adequacy of power distribution units, and configuration settings on relays as well as secondary unit substations. Because HPC data centers may be designed with little or no redundancy, a quality assurance plan is important for availability. For example, an independent third party can be used to verify torqued connections, quality of wiring, clearances, etc.

The power consumption of compute nodes in HPC data centers may vary greatly over short periods of time due to the intensive nature of the jobs. Peak power testing at Los Alamos National Laboratory using LINPACK shows voltage sags and swells up to 0.4% of RMS voltage with transients up to 1.5 MW [6]. Leibniz Supercomputing Centre has experienced power load fluctuations of 2 MW in time periods of much less than one second with their SuperMUC Phase 1 system running the SeisSol Benchmark [7]. Fig. 1 shows the very steep slopes and edges of the three-phase alternating current while executing the SeisSol application on 9216 nodes of SuperMUC Phase 1. The power draw of SuperMUC was in the order of 1 MW during the MPI-I/O phases of SeisSol and 3 MW during the highly optimized numerical simulation kernels of SeisSol.

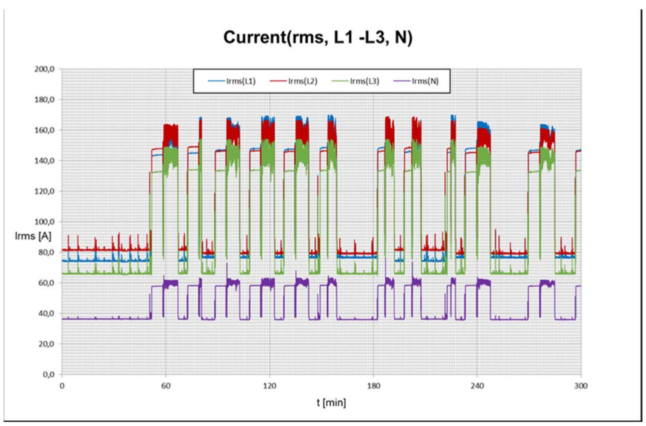

Fig. 1. Current During SeisSol Full System Run

With Exascale computing on the horizon, it is possible these short-term transients could be numerous MW in a fraction of a second and have an impact on the utility grid [8]. These types of transients should be checked during electrical commissioning.

Powerline harmonics is a concern due to the large concentrated quantity of IT equipment. While most IT equipment has power factor correction built into their power supply units (PSUs), there is the possibility of triplen harmonics that would appear on the neutral conductor of a 3-phase electrical distribution. Voltage distortion could also result in current harmonics. It is important to verify the electrical distribution can handle the harmonics and its heating effect on electrical distribution components.

To remove the heat load, liquid cooling is often used in the form of cold plates on processors and GPUs as well as rear door heat exchangers. The flow of cooling fluid requires electricity for pumps and valves that must meet the data center reliability, redundancy, and availability targets for uptime. The control sequences should be part of electrical commissioning to ensure cooling can be met under instantaneous power consumption changes.

Electrical commissioning must also be considered when it comes to power availability during migrations and de-installations. This is especially true when electrical power in a facility is at its limit and choices need to be made about installing more capacity or decommissioning equipment. This type of disruption should be checked carefully on paper to avoid the potential for unplanned downtime.

IV. Conclusion

Commissioning is not a one size fits all activity. Because data center design varies greatly, there are many technical differences and scenario specific considerations that will influence the commissioning, including overall schedule, timeline devoted to each phase of the project, cost, and the commissioning agent (e.g., phase when the agent is hired, experience). The Electrical System Commissioning Team under the leadership of the EE HPC WG needs to work together to develop a comprehensive guide to electrical commissioning that includes baseline requirements and additional requirements that account for the unique elements of HPC data centers.

[1] U.S. General Services Administration. "Site Commissioning White Paper." https://www.gsa.gov/cdnstatic/2017-13-07_SiteCommissioning_AccSingle_WhitePaper.pdf. July 2017.

[2] Building Commissioning Association. "New Construction Building Commissioning Best Practices." https://www.bcxa.org/wp-content/uploads/2016/03/BCxA.NCCx-BestPractices_031616.pdf. February 2016.

[3] Mills, E. "Building commissioning: a golden opportunity for reducing energy costs and greenhouse gas emissions in the United States." https://doi.org/10.1007/s12053-011-9116-8. 2011. Energy Efficiency 4, 145–173.

[4] Vallort, J., Weber, B. "Commissioning mission critical systems." Consulting-Specifying Engineer. 2009, 45 (5), pp. 34-39.

[5] Hydeman, M., Seidl, R., Shalley, C. "Data centers: Staying on-line: Data center commissioning." 2005. ASHRAE Journal, 47 (4), pp. 60-65.

[6] Backhaus S., Loncaric J., Wright C., Miller M. "HPC Power Transient Testing." 2013. Los Alamos National Laboratory, LA-UR-13-21919.

[7] A. Breuer, A. Heinecke, S. Rettenberger, M. Bader, A.-A. Gabriel und C. Pelties, Sustained Petascale Performance of Seismic Simulations with SeisSol on SuperMUC. In J.M. Kunkel, T. T. Ludwig und H.W. Meuer (Hrsg.), Supercomputing - 29th International Conference, ISC 2014, Band 8488 der Reihe Lecture Notes in Computer Science, S. 1–18. Springer, Heidelberg, Juni 2014. PRACE ISC Award 2014. doi: 10.1007/978-3-319-07518-1_1

[8] Huber, H. "Leibniz Supercomputing Centre." Presentation at 6th European workshop on HPC centre infrastructures, Såstaholm, Sweden. https://indico.fysik.su.se/event/4932/. 2015, May 11.