Applied Thermal Engineering 108 (2016) 937–950

Contents lists available at ScienceDirect

Applied Thermal Engineering

journal homepage: www.elsevier.com/locate/apthermeng

Research Paper

Analysis of airflow imbalances in an open compute high density storage data center

Husam A. Alissa [a,*], Kourosh Nemati [a], Udaya L.N. Puvvadi [b], Bahgat G. Sammakia, Fellow, IEEE [a], Ken Schneebeli [c], Mark Seymour [d], Tom Gregory [d]

[a] Department of Mechanical Engineering, ES2 Center, Binghamton University-SUNY, NY, USA
[b] Department of Computer Science, ES2 Center, Binghamton University-SUNY, NY, USA
[c] IBM Corporation, San Jose, CA, USA
[d] Future Facilities Ltd, London, UK

HIGHLIGHTS

- Experimental and numerical analysis of the open compute storage data center.
- Reliability and utilization impacts on IT during airflow imbalance in containment.
- Server and rack levels airflow prediction models.
- Airflow, pressure, and temperature flow curves to model IT fan control system.

ARTICLE INFO

Article history:
Received 7 May 2016
Revised 26 July 2016
Accepted 28 July 2016
Available online 29 July 2016

Keywords:
Storage
High density
IT
Containment
HDD
Open compute
Reliability
Flow curves
Utilization

ABSTRACT

This article reports experimental and numerical testing performed to characterize the operation and reliability of the open compute (OC) storage system in contained environment from server to aisle levels. The study is comprised of three parts. The first part is an experimental analysis of the high density (HD) 3D array storage unit thermal and utilization responses during airflow imbalances. This is done with the stress test proposed for IT in containment to mimic possible mismatch and cascade failure scenarios. It is found that downstream HDDs are most prone to overheating and loss in utilization during an airflow imbalance. This was proven to undermine the storage capacity of the hard disk drives. An IT level airflow prediction model is discussed for the storage unit and validated for different fan speeds. In the second part, a computational fluid dynamics model is created for a high density open rack based on the active flow curve method. Here, the measured airflow response curves for the open compute IT (storage and compute servers) are used to build compact models and run rack level testing for IT air systems sensitivity and create a rack level AFC (active flow curve) airflow demand prediction model. Finally, the experimental characterization data is used to build an aisle level model (POD) that incorporates IT fan control systems (FCS). This modeling approach yields shorter uptime during chiller failure due to increased recirculation induced by increased IT airflow demand during cases such as chiller failure or high economizer temperatures.

1. Introduction

Every data hall (data center) is expected to deliver information technology services such as telecom, compute, storage or usually a mix of all three. IT equipment can be divided into three main categories based on their functions: 1 – Compute/Web Servers; 2 – Network Switches; 3 – Storage units (servers). Each category IT equipment has three major requirements to guarantee its functionality: space, power and sufficient cooling. In terms of cooling, IT equipment research has focused on the compute category of IT. This is due to the high frequency central processing unit (CPU), which forms the building block of computing data center infrastructure. This particular focus is justifiable given Moore's law, which relates the increased transistor density with the increase in heat fluxes.

* Corresponding author.
E-mail address: halissa1@binghamton.edu (H.A. Alissa).

http://dx.doi.org/10.1016/j.applthermaleng.2016.07.186
1359-4311/

Nomenclature

A	proportional multiplier or area
AF	acceleration factor
AFC	active flow curve
ARFS	average rack fans' speed
ASIC	application-specific integrated circuit
CAC	cold aisle containment
C_p	specific heat at constant pressure
CPU	central processing unit
D_h	hydraulic diameter (m)
DPB	drive plane board
DRAM	dynamic access random memory
Ea	activation energy (eV)
EB	exabyte (1 billion Gigabytes)
FCB	fan control board
FCS	fan control system
FD	free delivery (CFM)
F_i	body force in Cartesian direction i
HAC	hot aisle containment
HBA	hot bus adapter
HD	high density
HDD	hard disk drive
IPMI	Intelligent Platform Management Interface
IT	information technology, (also refers to servers, storages, and switches)
K	Boltzmann's constant (eV/K)
K	effective thermal conductivity
JBOD	just a bunch of disks
$K_{i,v}$	modified inertial (in. H_2O/CFM^2) and viscous (in. H_2O/CFM) loss coefficients
$K_{1,2}$	inertial (in. H_2O/CFM^2) and viscous (in. H_2O/CFM) loss coefficients
μ	dynamic viscosity (kg/m s)
N	fan speed
NFS	network file system
Nu	Nusselt Number
OC	open compute
OR	open rack
OU	OpenU, rack unit height, 1 OU = 1.89″ (48 mm)
P	pressure (in. H_2O)
PC	personal computer
P_C	critical pressure (in. H_2O)
PCB	printed circuit board
PCIe	peripheral component interconnect express
PFC	passive flow curve
PI	partial integral controller
POD	performance optimized datacenter
PXE	preboot execution environment
Q	airflow (CFM)
Q_o	operational airflow (CFM)
$Q_\%$	percentage airflow reduction from FD
ρ	density (kg/m^3)
Re	Reynolds number
SEB	SAS expansion board
S_h	energy source per unit volume
SMART	data from a hard drive or solid state drive's self-monitoring capability
SSD	solid state drive
T	temperature
$T_{2,1}$	absolute temperatures during stressed and normal operations
τ_{ij}	stress tensor
u_i	velocity component in Cartesian direction i
V	average velocity (m/s)
VFD	variable frequency drive

Growth in data storage is being fueled by social media usage and the large volume of images and videos that are being generated from mobile devices, and personal computers (PC). At the same time, machines themselves are generating more data than ever before. An example is the Large Hadron Collider that can generate about 40 TB of data per second during a simulation of the Big Bang model. Note that even a user's sixty seconds online transaction with a retailer can generate a significant amount of data on the products, prices, payment methods, manufacturer, and credit history involved. Consequently, it is no surprise that worldwide data storage has increased by nearly 8000 Exabytes (EB) from 2005 to 2015 [1].

Data storage can be accomplished with Solid State Drives (SSD), Mechanical Hard Disk Drives (HDD), or Tape. SSDs which do not contain any mechanical parts, store data on interconnected flash memory chips. Mechanical HDDs utilize rotationally spinning magnetic media for data storage. The main difference with SSD is that HDD technology is more mature and of significant lower cost per unit storage which also fits perfectly in applications such as cold storage or archive. For HDDs, the power generated in the disk drives arises from mechanical frictional resistance, motor work, and from electronic package heat dissipation. Its primary components are the aluminum die casting (housing), the sheet metal cover, the spindle attached to a motor and several disk shaped platters, and a printed circuit board (PCB) with memory and an application-specific integrated circuit (ASIC) devices to allow for interface and communication between the magnetic media and the external devices and systems. The case and mechanical parts of the HDD are shown Fig. 1(a), the arm motion is derived through the Lorentz force by passing a current through the coil in the presence of a magnetic field.

Storage servers often place multiple HDDs in close proximity to each other with narrow passages or air gaps between them to facilitate airflow. Historically, these storage server designs place the HDDs at the front of the server so as to allow users to insert and eject the drives with minimal impact to the server's operation during drive replacement. In the pursuit of higher volumetric drive densities in the rack and Open Compute (OC) designs, there are newer server products which incorporate denser configurations in which the hard disk drives fill up most of the volume of the storage server, not just the front access area – as in the legacy storage servers. In such design, it is important to understand various thermo-fluid effects including the temperature profile of the hard disk drive, interactions between the drive and airflow, and the relationship between the air gap and the system level pressure drop. Note that there is very rare published data on the aerodynamic and thermal interactions of storage unit chassis and the HDDs to different environmental conditions let alone high density (HD) OC storage servers.

The Open Compute Project (OCP) aims to design and share intellectual property on efficient servers and storages (data center hardware in general) for scalable computing. The thermal design of an OC data center is discussed in [2]. The study explained Facebook design of an OC 100% outside air economization facility that uses evaporative cooling and hot aisle containment (HAC). The storage units are usually accompanied by a web server to manage the storage. Optimizing performance through modeling and measurements for an open compute web server unit is discussed