SNIA SFF TWG Technology Affiliate

SFF-TA-1023

Specification for

Thermal Characterization Specification for EDSFF Devices

Rev 0.8.5 August 12, 2021

SECRETARIAT: SFF TA TWG

This specification is made available for public review at http://www.snia.org/sff/specifications. Comments may be submitted at http://www.snia.org/feedback. Comments received will be considered for inclusion in future revisions of this specification.

The description of the connector in this specification does not assure that the specific component is available from connector suppliers. If such a connector is supplied, it should comply with this specification to achieve interoperability between suppliers.

ABSTRACT: This specification defines thermal performance measurement fixtures, methods, and reporting requirements for Enterprise/Datacenter Small Form Factor (EDSFF) devices. EDSFF device thermal characteristics are influenced by the temperature and airflow velocity, enclosure design, workload and power requirements. This specification establishes a common method to characterize the variety of EDSFF devices such that system integrators are enabled to evaluate whether a given device can be adequately cooled based on the system's dynamic operating capabilities.

POINTS OF CONTACT:

William Lynn
Senior Distinguished Engineer
Dell EMC
One Dell Way
Round Rock, TX 78683
Email: William.Lynn@Dell.com

Pranay Mahendra
Master Technologist
Hewlett-Packard Enterprise
11445 Compaq Center Drive
Houston, TX 77070-1433
Email: pranay.mahendra@hpe.com

Chairman SFF TA TWG
Email: SFF-Chair@snia.org

Intellectual Property

The user's attention is called to the possibility that implementation of this specification may require the use of an invention covered by patent rights. By distribution of this specification, no position is taken with respect to the validity of a claim or claims or of any patent rights in connection therewith.

This specification is considered SNIA Architecture and is covered by the SNIA IP Policy and as a result goes through a request for disclosure when it is published. Additional information can be found at the following locations:

- Results of IP Disclosures: http://www.snia.org/sffdisclosures
- SNIA IP Policy: http://www.snia.org/ippolicy

Disclaimer

The information contained in this publication is subject to change without notice. The SNIA makes no warranty of any kind with regard to this specification, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. The SNIA shall not be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance, or use of this specification.

Suggestions for revisions should be directed to http://www.snia.org/feedback/.

Foreword

The development work on this specification was done by the SNIA SFF TA TWG, an industry group. Since its formation as the SFF Committee in August 1990, the membership has included a mix of companies which are leaders across the industry.

For those who wish to participate in the activities of the SFF TA TWG, the signup for membership can be found at http://www.snia.org/sff/join.

Revision History

Rev 0.1 *August 13, 2020*
- Initial Draft

Rev 0.2 *August 13, 2020*

Rev 0.3 *January 25, 2021*

Rev 0.4 *February 2, 2021*

Rev 0.5 *February 3, 2021*
- Added verbiage around Dtherm levels

Rev 0.6 *February 4, 2021*
- Added developers note around intended use for Dtherm levels

Rev 0.7 *February 26, 2021*
- Refined recommended design space, added UUT definition

Rev 0.8 *March 17, 2021*
- Refined MaxTherm/Dtherm measurement requirements, specifically allowing for analytical temperature scaling based on device thermal margin
- Clarified that Dtherm levels should correlate to NVMe power states for E3 NVMe use cases
- Clarified that in the 6xE3 thin test case, only 3 adjacent devices need to be functional and stressed for thermal measurement
- Updated the E3.S 1T recommended design space

Rev 0.8.1 *April 5, 2021*
- Corrected Table 4-3 to show Y axis as Channel Velocity (LFM)
- Updated Figure 4-1 to represent the larger range indicated by the equations in Table 4-1
- Added a note that AFI static pressure and airflow measurements should be corrected to sea level altitude and 25C ambient temperature

Rev 0.8.2 *April 6, 2021*
- Added table 5-1, which provides a template for device makers to show device power and performance at Maxtherm and Dtherm levels, as well as other parameters relevant to E3 thermal performance
- Added a detailed description and equation in Section 5 for analytical temperature scaling in the test fixture during Maxtherm/Dtherm testing

Rev 0.8.3 *May 13, 2021*
- Added drawings for E1 test fixtures, updated AFI levels and Thermal Levels to include other EDSFF Form Factors
- Reformatted specification to reflect addition of E1
- Added HT–LF design space (high temperature, low flowrate), and isometric views of E1 test fixture
- Added MinAmbient definition

Rev 0.8.4 *July 9, 2021*
- Corrected typos
- Added section 3.4 (Units and Unit conversions)
- Added figure in section 4.1 showing example usage of AFI level measurement
- Added section 4.5, which details the HT-LF design space
- Added SFF-TA-1007 as an additional reference document

Rev 0.8.5 *August 13, 2021*

- Changed the E1.L pitch from 12.5mm to 13mm to align with E1.S in Figure 5-10
- Removed the reference to REF-TA-1011 in section 2.1
- Added ISO as a source for documents in section 2.2
- Corrected the definition of **Restricted:** in section 3.1
- Clarified Developer's Note in section 4.3
- Rewrote the introduction to the second paragraph in section 4.4
- Added ISO 5081:2017 as a heat chamber option in section 5
- Harmonized labels between figures 5-2, 5-3, and 5-4
- Updated Figures 5-5 and 5-10
- Add a developer's note below Figure 5-2 with recommended material and device orientation

Contents

Figures

Tables

1. Scope

This specification defines a methodology for thermal performance characterization of EDSFF devices defined within SFF-TA-1008 (E3) and SFF-TA-1006 (E1). This specification defines the test methodology, simulation methodology, test fixtures, and parameters to be reported for thermal characterization.

2. References and Conventions

2.1. Industry Documents

The following documents are relevant to this specification:

- ASME Y14.5-2009	Dimensioning and Tolerancing
- ANSI/AMCA210-07, ISO 5801:2017	Performance testing using standardized airways
- SFF-TA-1002	Protocol Agnostic Multi-Lane High Speed Connector
- SFF-TA-1006	Enterprise and Datacenter 1U Short SSD Form Factor (E1.S)
- SFF-TA-1007	Enterprise and Datacenter 1U Long SSD Form Factor (E1.L)
- SFF-TA-1008	Enterprise and Datacenter SSD 3" Form Factor (E3)

2.2. Sources

The complete list of SFF documents which have been published, are currently being worked on, or that have been expired by the SFF Committee can be found at http://www.snia.org/sff/specifications. Suggestions for improvement of this specification will be welcome, they should be submitted to http://www.snia.org/feedback.

Copies of ASME standards may be obtained from the American Society of Mechanical Engineers (https://www.asme.org).

Copies of Electronic Industries Alliance (EIA) standards may be obtained from the Electronic Components Industry Association (ECIA) (https://www.ecianow.org).

Copies of American National Standards Institute Documents may be obtained from the American National Standards Institute (ANSI) at https://www.ansi.org/

Copies of ISO documents may be obtained from the International Organization for Standards at https://www.iso.org

2.3. Conventions

The following conventions are used throughout this document:

DEFINITIONS

Certain words and terms used in this standard have a specific meaning beyond the normal English meaning. These words and terms are defined either in the definitions or in the text where they first appear.

ORDER OF PRECEDENCE

If a conflict arises between text, tables, or figures, the order of precedence to resolve the conflicts is text; then tables; and finally figures. Not all tables or figures are fully described in the text. Tables show data format and values.

LISTS

Lists sequenced by lowercase or uppercase letters show no ordering relationship between the listed items.

EXAMPLE 1 - The following list shows no relationship between the named items:

a. red (i.e., one of the following colors):
    A. crimson; or
    B. pink;
b. blue; or
c. green.

Lists sequenced by numbers show an ordering relationship between the listed items.

EXAMPLE 2 -The following list shows an ordered relationship between the named items:

1. top;
2. middle; and
3. bottom.

Lists are associated with an introductory paragraph or phrase, and are numbered relative to that paragraph or phrase (i.e., all lists begin with an a. or 1. entry).

DIMENSIONING CONVENTIONS

The dimensioning conventions are described in ASME-Y14.5-2009, Geometric Dimensioning and Tolerancing. All dimensions are in millimeters, which are the controlling dimensional units (if inches are supplied, they are for guidance only).

NUMBERING CONVENTIONS

The ISO convention of numbering is used (i.e., the thousands and higher multiples are separated by a space and a period is used as the decimal point). This is equivalent to the English/American convention of a comma and a period.

American	French	ISO
0.6	0,6	0.6
1,000	1 000	1 000
1,323,462.9	1 323 462,9	1 323 462.9

3. Keywords, Acronyms, and Definitions

For the purposes of this document, the following keywords, acronyms, and definitions apply.

3.1. Keywords

May / may not: Indicates flexibility of choice with no implied preference.

Obsolete: Indicates that an item was defined in prior specifications but has been removed from this specification.

Optional: Describes features which are not required by the SFF specification. However, if any feature defined by the SFF specification is implemented, it shall be done in the same way as defined by the specification. Describing a feature as optional in the text is done to assist the reader.

Prohibited: Describes a feature, function, or coded value that is defined in a referenced specification to which this SFF specification makes a reference, where the use of said feature, function, or coded value is not allowed for implementations of this specification.

Reserved: Defines the signal on a connector contact [when] its actual function is set aside for future standardization. It is not available for vendor specific use. Where this term is used for bits, bytes, fields, and code values; the bits, bytes, fields, and code values are set aside for future standardization. The default value shall be zero. The originator is required to define a Reserved field or bit as zero, but the receiver should not check Reserved fields or bits for zero.

Restricted: Refers to features, bits, bytes, words, and fields that are set aside for other standardization purposes. If the context of the specification applies the restricted designation, then the restricted bit, byte, word, or field shall be treated as a value whose definition is not in scope of this document, and is not interpreted by this specification.

Shall: Indicates a mandatory requirement. Designers are required to implement all such mandatory requirements to ensure interoperability with other products that conform to this specification.

Should: Indicates flexibility of choice with a strongly preferred alternative.

Vendor specific: Indicates something (e.g., a bit, field, code value) that is not defined by this specification. Specification of the referenced item is determined by the manufacturer and may be used differently in various implementations.

3.2. Acronyms and Abbreviations

E1: Form factors and devices defined in SFF-TA-1006 and SFF-TA-1007
E3: Form factors and devices defined in SFF-TA-1008
EDSFF: Enterprise/Datacenter Small Form Factor

3.3. Definitions

AFI Level – Air Flow Impedance Level – used to compare airflow impedance between devices of the same form factor

MaxTherm Level - Minimum airflow required at a given approach air temperature for which a device operating at TDP will operate without degraded performance. TDP refers to 'Thermal Design Power', which is the maximum amount of heat capable of being generated by the device under load.

DTherm Level - Minimum airflow required at a given approach air temperature for which a device will operate but at a reduced device performance level. More than one Dtherm level can be set for a device under different performance modes. Dtherm levels should coincide with NVMe power states for NVMe EDSFF devices.

MaxAmbient - Approach air temperature upper threshold if required by a device (for example, some devices may not be designed to operate above 60 °C)

MinAmbient – Approach air temperature lower threshold if required by a device (for example, some devices may not be designed to operate below 0 °C)

Approach Air Temperature – Average temperature of the air before it reaches the EDSFF devices

Channel Velocity –Velocity of air when traveling in the gaps between the EDSFF devices. See Section 5.

3.4. Units and Unit Conversions

Parameter	Standard Units	SI	Conversion
Volumetric Flowrate	CFM (Cubic Feet per Minute)	M^3/S (Cubic Meters per Second)	1 CFM = 4. 72 E-4 M^3/s
Velocity	LFM (Linear Feet per Minute)	M/s (Meters per Second)	1 LFM = 5. 08 E-3 M/s
Pressure	inH$_2$0 (Inches of Water)	Pa (Pascals)	1 inH$_2$0 = 248. 84 Pa
Length	In (inches), Ft (feet)	mm (millimeter)	1 foot = 12 inches = 304. 8mm
Temperature	Fahrenheit (˚F)	Celsius (˚C)	(32˚F -32) * 5/9 = 0˚C

Table 3-1: Units and Unit Conversions

4. EDSFF Device Thermal Reporting

4.1. Overview

This document defines thermal characterization and parameter reporting for EDSFF devices and describes the test and modelling procedure for the measurement of thermal and airflow impedance of an array of devices. Using this framework, specific implementations of EDSFF devices may be characterized and their parameters communicated effectively to enable system integrations.

Thermal Reporting uses three metrics that characterize the device airflow impedance levels and cooling requirements. These metrics are AFI, MaxTherm, and DTherm. The metrics allow devices to be classified into groups so that system designers may select devices that are compatible with their host systems. In addition, system software may retrieve the metrics from the device and utilize that information to set appropriate fan speeds. Device designers are encouraged to utilize the metrics and work with system designers to ensure that the devices being developed, and their associated metrics are aligned with system capabilities.

4.2. Airflow Impedance (AFI) Level

The intent of Airflow Impedance Level information is to classify devices into groups based on their impedance to airflow. These groups are referred to as AFI levels for the purpose of this specification. The AFI level is an important parameter for system designers due to the impact to system airflow that a high airflow impedance device may have versus a low airflow impedance device. The AFI level is important for system cooling as well as individual device cooling. The installation of higher airflow impedance devices, especially if multiple such devices are installed, in a platform may exceed the capabilities for a given platform fan performance curve. Through testing, system developers determine which device AFI levels a platform can support. The AFI level may also be used in real time by system fan control algorithms to manipulate fan speeds based on the AFI reported by the devices.

Section 5 details the test procedure for determining a device's AFI level and provides the appropriate test fixture geometry for each type of EDSFF device. Device airflow impedance is categorized as one of 8 levels communicated by the AFI level field. Each AFI level corresponds to one of the curves illustrated in Figure 4-1. The AFI Level field is reported as a value of one through eight.

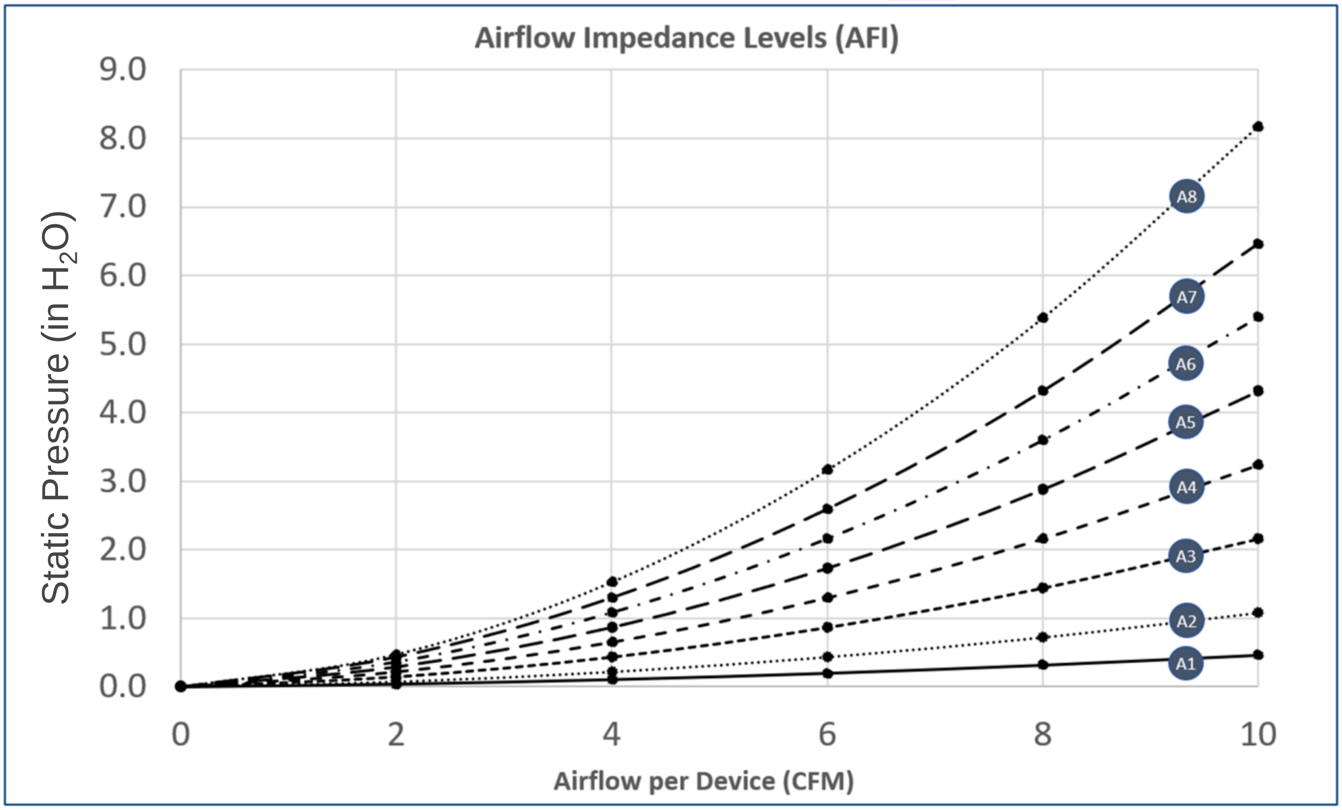

Figure 4-1: AFI Levels for EDSFF Devices

A device is assigned to an AFI level by comparing its measured airflow impedance in the appropriate test fixture within the range of 2 to 10 CFM per device to the AFI Level equations listed in Table 4-1. The device is assigned the highest of the AFI levels intersected by the device airflow impedance curve. The airflow vs static pressure measurements used to characterize a devices AFI level should be corrected to equivalent values at sea level and 25 °C if the test is conducted at a different altitude or ambient temperature. An example of proper AFI characterization is displayed below in Figure 4-2. Here, UUT (unit under test) represents a hypothetical device whose airflow impedance is characterized in the appropriate test fixture. In this instance, the UUT would be assigned AFI level 2, as that is the highest AFI level that the device's airflow impedance curve intersects.

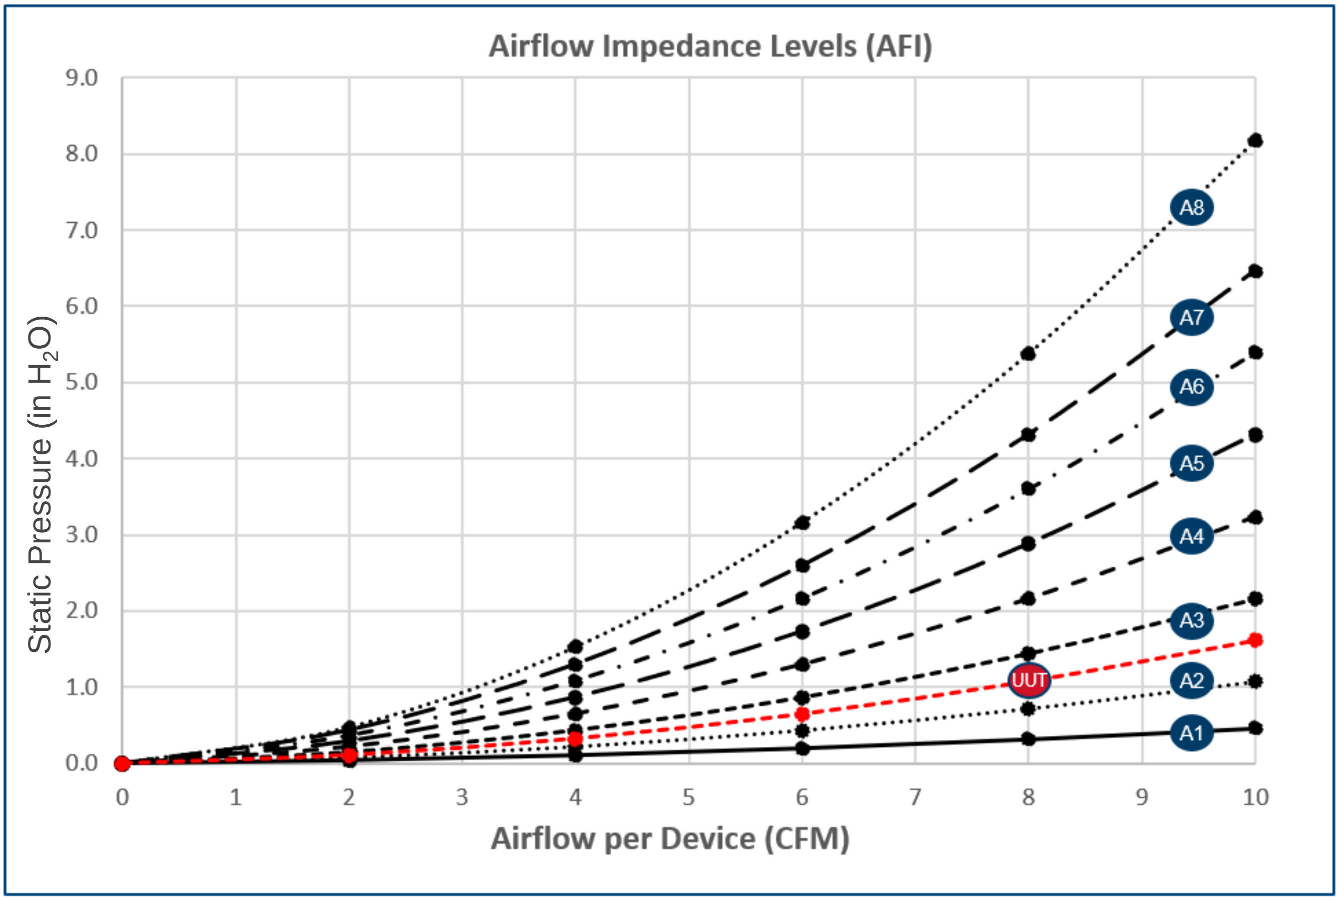

Figure 4-2: Example Device AFI Level

Table 4-1: AFI Level Equations

AFI Level	Equation (X=volumetric airflow [CFM]; Y=static pressure [in. H20])
0	Reserved
A1	Y = 0. 003 30 * X^2 + 0. 013 30 * X
A2	Y = 0. 008 93 * X^2 + 0. 018 57 * X
A3	Y = 0. 017 85 * X^2 + 0. 037 13 * X
A4	Y = 0. 026 78 * X^2 + 0. 055 70 * X
A5	Y = 0. 035 70 * X^2 + 0. 074 27 * X
A6	Y = 0. 044 63 * X^2 + 0. 092 83 * X
A7	Y = 0. 053 55 * X^2 + 0. 111 40 * X
A8	Y = 0. 072 60 * X^2 + 0. 091 50 * X

Developer's Note: *It is imperative that the appropriate test fixture (as defined in Section 5) is used to determine AFI level for a given device. Due to differences in test fixture geometry, it is not accurate to directly compare airflow impedance of different form factors. AFI level is a useful metric for comparing two devices of the same form factor but is not to be used to compare different form factor devices.*

4.3. Maximum Thermal (MaxTherm) Level

The intent of the Maximum Thermal (MaxTherm) Level information is to define the minimum airflow rate required at a given air temperature for which a device, when stressed to its TDP (thermal design power) level limit, will operate within its component's reliability limits (as defined by the device/part manufacturer) and without degraded device performance. Once a device's thermal performance profile is established, the MaxTherm Level is determined. MaxTherm is the lowest curve in Figure 4-3 which is entirely above the device's quantified thermal performance curve throughout the range of approach ambient temperatures (also referred to as approach temperature) illustrated on the X-axis ranging from 25 °C up to 65 °C.

Airflow and thermal measurement shall be performed in the appropriate thermal test fixture for the device being tested, as defined in Section 5. The general process to determine the MaxTherm Level is as follows:

1. Setup devices in the appropriate test fixture with cables or backplane connecting devices to host system or controller (keep-in volumes for cables in the test fixtures are defined in Section 5)
2. Set flow bench or wind tunnel to appropriate airflow setting for desired measurement point
3. Initiate device stress required to achieve the maximum device TDP
4. Allow device temperatures to reach steady state
5. Collect device temperatures in addition to local approach temperature and airflow measurements
6. Approach temperature and measured device temperatures are scaled up until the UUT (unit under test) reaches its maximum component temperature. This approach temperature and air flow rate (per device) pair is one point on the device cooling curve. This temperature scaling can be performed experimentally via heating the approach air, or analytically by holding approach air temperature constant and using remaining thermal margin to determine the maximum allowable approach temperature for a given flowrate. For a detailed explanation of the analytical approach for air temperature scaling, see Section 5.
7. Repeat for other airflow rates as defined by Figure 4-3 to develop the complete device cooling curve
8. The device cooling curve (required CFM/device vs. local approach temperature) is plotted against the predefined MaxTherm levels shown in Figure 4-3. The device MaxTherm level is the lowest curve in the figure which is entirely above the device cooling curve throughout the range of approach temperatures illustrated on the X-axis ranging from 25 °C up to 65 °C. An example of this is illustrated in Figure 4-4, where the thermal performance of a hypothetical UUT (unit under test) is displayed. In this instance, the UUT would be classified as being a thermal level 2 device in the tested operating state.

The thermal levels defined in Figure 4-3 display CFM/dev on the Y-axis, and approach air temperature on the X-axis.

Developer's Note: *Airflow is recorded as CFM/device, but channel velocity can be calculated easily for uniformly shaped devices using Equation 5-1 as defined in Section 5. An example of this is shown in Figure 4-3. Note that the recommended design space for EDSFF devices is shown in the blue shaded region on the chart. For EDSFF devices, the recommended design space is below Level 3, up to 60 °C approach velocity. A high-temperature, low airflow design space is specified as well, and detailed in 4.5*

- *Curves in Figure 4-3 are shown along with the recommended design space.*
- *Compare measured readings against the equations defined in Table 4-2.*
- *Thermal Levels 3-8 do not require a device to support less than 1.5 CFM.*
- *Not every device is expected to perform within the **HT-LF** space – this area is for targeted device designs or reduced performance applications as defined in 4.5.*

Developer's Note: *It is strongly recommended that device developers strive to offer solutions that can operate within the shaded recommended design space envelope illustrated in Figure 4-3 to maximize the number of systems that can support such conditions. Some devices may not meet the recommended design space, e.g., devices consuming 70 W. Systems intended to incorporate such devices have practical implications to implement thermal mitigations.*

The MaxTherm field is defined for a value of 1 through 8.

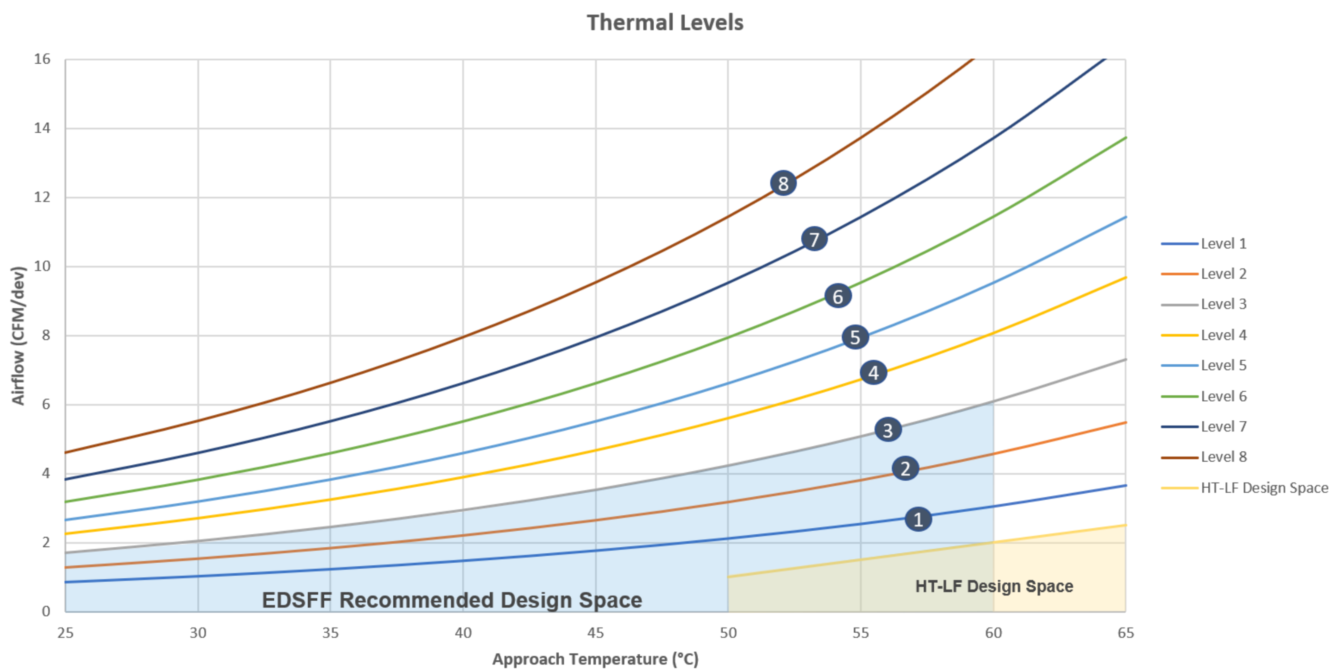

Figure 4-3: MaxTherm and DTherm Levels

No lower flow rates shall be defined for approach air temperatures below 25 °C.

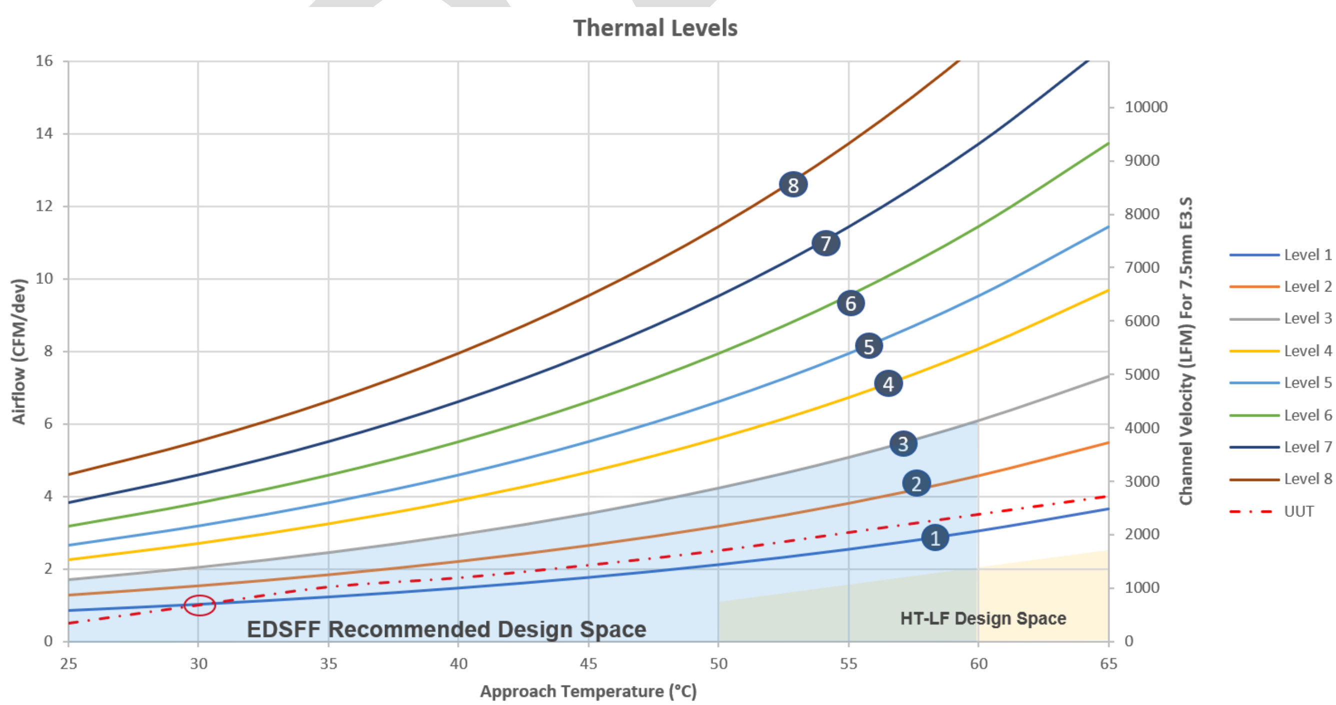

Figure 4-4: Example Device Thermal Profile

As an example of proper EDSFF test fixture usage, Figure 4-4 includes the curve of a hypothetical device that has been tested. The highlighted circle illustrates that the UUT (Unit Under Test) exceeds MaxTherm Level 1 above an approach air temperature of 30 °C. Given this example device cooling curve, MaxTherm Level 2 would be selected for this device when running at its maximum TDP. The UUT in this instance is assumed to be a 7.5mm E3 1T device, and channel velocity has been calculated using Equation 5-1 and is displayed on the right-hand Y-axis. Table 4-2 provides the Thermal Level equations as displayed in Figure 4-3.

Table 4-2: Thermal Level Equations (CFM/Device)

Thermal Level	X = Approach Temperature; Y = Airflow Required (CFM/Device)
0	Reserved
1	Y = 0. 341 6 * e^(0. 036 5 * X)
2	Y = 0. 521 4 * e^(0. 036 5 * X)
3*	Y = 0. 683 2 * e^(0. 036 5 * X)
4	Y = 0. 904 2 * e^(0. 036 5 * X)
5	Y = 1. 068 9 * e^(0. 036 5 * X)
6	Y = 1. 282 8 * e^(0. 036 5 * X)
7	Y = 1. 539 4 * e^(0. 036 5 * X)
8	Y = 1. 847 2 * e^(0. 036 5 * X)
HT-LF**	Y = 0.1 * X -4

Note* *Level 3 is the recommended design space limit for EDSFF devices.*

Note** *HT-LF (High Temp, Low Flow) is a specific design space not expected to be targeted for every device.*

4.4. Designed Thermal (DTherm) Level

The intent of the Designed Thermal (DTherm) Level information is to determine the minimum airflow required at a given air temperature for which a device will operate within its component's reliability limits but at a reduced device performance level. The stress application that was used for MaxTherm is also used for DTherm. Typically, the reduced performance is induced by the device's self-initiated thermal protection schemes, such as throttling. This is of interest for platforms unable to provide sufficient cooling, such as due to the loss of a cooling fan which reduces the platform's cooling capacity. The reduced cooling capacity in such a case may be below the device's MaxTherm Level.

If a device supports multiple designed power performance levels and these translate into reduced thermal profiles than established for MaxTherm, then DTherm is not equal to MaxTherm. DTherm is calculated using the same process as MaxTherm, however, the device's minimum Dtherm profile is created using the device's most reduced performance operating level at which it is still operating and providing useful work as defined by the device implementer. The number of reported DTherm levels are up to the device implementer based on software definition and numbers of operating states. DTherm reporting is strongly recommended as this information allows system integrators to understand performance throttling options to balance performance with available system cooling. Devices may be set to run continuously in a DTherm level or be allowed to self-regulate Dtherm levels based on available cooling. DTherm levels should be used to represent a device in different power states as defined by the appropriate industry standard for that device type (i.e., NVMe devices would utilize the NVMe Power States). Additionally, DTherm levels can be used to enable a device to perform within the recommended EDSFF design space or the HT-LF design space.

Developer's Note: *It is strongly recommended that device developers strive to achieve better thermal performance by establishing multiple DTherm levels. Not all slots in all systems will be optimally located for EDSFF device cooling; creating a means for a device to function (even in a reduced performance state) in a wide variety of boundary conditions ensures that the device can be used broadly across a system portfolio. For instance, a device will have a MaxTherm level representing thermal performance at TDP and maximum bandwidth but should also have multiple DTherm levels beneath this showing improved thermal performance at reduced bandwidth and device power. The device should be able to switch to a DTherm level of performance automatically in the event of a* *thermal excursion but should also contain a means by which the controller can initiate a DTherm level limit for the device. An example of proper Dtherm usage would be as follows: MaxTherm represents 100% bandwidth, DTherm 1 represents 75%, DTherm 2 represents 50%, DTherm 3 represents 25%, DTherm 4 represents maximum throttled condition while still providing useful performance. Additional DTherm levels may be used to add further granularity as needed.*

4.5. High Temperature – Low Flow Design Space (HT-LF)

In addition to the recommended design space, some applications will require High Temperature, Low Flowrate operating conditions (**HT-LF**). It is not expected that a device will have to achieve maximum performance within the HT-LF space, but reliability and performance capability should be understood and reported for devices designed to operate within this space. The HT-LF space is defined by the yellow shaded region in Figure 4-3. Typical use cases for HT-LF devices might include rear storage slots, or for environments where system inlet temperature is expected to be 50˚C or greater during normal operation.

4.6. MaxAmbient

Some devices may not be thermally viable up to an approach air temperature of 65 °C, regardless of the air speed involved. Or, even if the device could be thermally viable at high temperature extremes, the airflow requirements to do so may require it to operate at a higher thermal challenge level, potentially reducing the number of systems able to supply the needed airflow. In either circumstance, the device may opt to have its upper approach air temperature limit be less than 65 °C. This approach air temperature upper threshold is referred to as MaxAmbient and is defined as the temperature just below or at a throttling condition. MaxAmbient shall be an integer value between 50 °C and 65 °C inclusive. Max Ambient applies to the full operating potential MaxTherm condition.

4.7. MinAmbient

Like MaxAmbient, some devices may not be thermally viable below a given air temperature. The thermal levels defined in Figure 4-3 do not assume any ambient temperatures lower than 25 °C. In some operating environments, ambient temperature may fall below 25 °C, and therefore a lower threshold for ambient temperature should be established by the device manufacturer. The MaxTherm level of a device at a temperature below 25 °C can be determined based on that device's performance from 25 °C to MaxAmbient.

5. Thermal Data Collection and Test Procedure

The following specifies the thermal test set-up and procedure to evaluate an EDSFF device's thermal performance. As previously discussed, thermal performance is partitioned into five sub-fields:

- AFI Level – Air Flow Impedance Curve
- MaxTherm Level - Minimum airflow required at a given air temperature for which a device at TDP level limit will operate and without degraded performance
- DTherm Level - Minimum airflow required at a given air temperature for which a device, when provided the same stress application as it was for MaxTherm, will operate but at a reduced device performance level
- MaxAmbient - Approach air temperature upper threshold between 50 °C and 65 °C if required by a device
- MinAmbient – Approach air temperature lower threshold at some value below 25 °C if required by a device.

These five sub fields should be used to populate Table 5-1 for the unit under test:

Table 5-1: Device Thermal Performance Template

Device Form Factor:				
Device AFI Level				
Device Max Ambient				
Device Min Ambient				
HT-LF Capable (Y/N)				
	Device Power (W)	**Thermal Level**	**Seq RD (MB/s)**	**Seq Wr (MB/s)**
MaxTherm				
DTherm 1				
DTherm 2				
DTherm 3				
DTherm 4				
DTherm 5 (as needed)				

In the above table, performance should be specified in Megabytes per second (MB/s). The values of these fields are determined using the test fixture. Details regarding the dimensions and construction of the test fixture are provided in this specification along with the associated CAD files (contained in the SFF-TA-1023.zip file from SNIA SFF TA TWG). The test fixture is intended to be attached to an AMCA 210-99/ASHRAE 51-1999/ISO 5081:2017 compliant airflow chamber, an example of which is shown in Figure 5-1, which can quantify both static pressure as well as volumetric airflow.

Figure 5-1: Airflow Chamber

The tester is attached to the airflow chamber such that ambient air is pulled into the fixture from the LED facing side of the device. When performing airflow impedance tests to determine the AFI level only the devices should be installed into the fixture. No backplane or cables should be included. Airflow set points for AFI levels are 2, 4, 6, 8, and 10 CFM per device. Thus, for an example of a 3 device fixture the airflow points are 6, 12, 18, 24, and 30 CFM. An example of a fixture with 6 devices will have airflow points of 12, 24, 36, 48, and 60 CFM. At each airflow point the static pressure at the flow bench is collected in units of inches of water (in. H_20.).

Additionally, Channel Velocity or Approach Velocity can be calculated to aid in system design. As device pitch and thickness can vary from system to system, channel velocity can be used in place of CFM/device for regularly shaped devices without extended fin surfaces. Channel velocity can be calculated easily as follows using Equation 5-1 (which contains the unit conversions necessary to input CFM/dev and device thickness/pitch in millimeters):

$$\textbf{Equation 5-1: } Channel\ Velocity\ (LFM) = \frac{\frac{CFM}{Dev} * 92\ 903}{(Drive\ Pitch\ (mm) - Drive\ Thickness\ (mm)) * Drive\ Width\ (mm)}$$

For example, channel velocity in the test fixture for a 7.5mm E3 1T device at 2.5 CFM/dev would be found as follows:

$$\textbf{Example: } Channel\ Velocity\ (LFM) = \frac{2.5\ CFM/dev * 92\ 903}{(9.3mm - 7.5mm) * 76mm} = \mathbf{1\ 697}\ LFM$$

Here, channel velocity is defined as the gap between two uniformly shaped EDSFF devices. The thermal levels defined in Figure 4-3 display CFM/dev on the left-hand Y-axis, and channel velocity (in the test fixture) on the right-hand Y-axis. Calculating channel velocity may allow more accurate predictions of device performance in system slots, as device pitch and device thickness can vary. Additionally, approach velocity can also be derived from CFM/dev, using Equation 5-2 below:

$$\textbf{Equation 5-2: } Approach\ Velocity\ (LFM) = \frac{\frac{CFM}{Dev} * 92\ 903}{(Drive\ Pitch\ (mm)) * Drive\ Width\ (mm)}$$

Both Channel Velocity and Approach Velocity are derived from CFM/dev, which is the recommended unit for evaluating device thermal performance in a test fixture. Channel Velocity and Approach Velocity may be used by system designers as needed, but only CFM/dev should be reported by the device designer.

When testing to determine the MaxTherm or DTherm levels the test fixture and device to be tested must be setup according to this specification. As with the AFI level testing, the tester is attached to the airflow chamber such that ambient air is pulled into the fixture from the LED facing side of the device. Identical devices must be installed in the slots adjacent to the device under test. All cables needed to fully exercise the three devices are attached to the devices and routed outside the system to minimize the blockage of airflow. While the E3 1T test fixture can hold up to 6 devices, only 3 adjacent devices need to be stressed for each measurement. The 3 devices not stressed need to share the same geometry with those being stressed so that airflow is uniform through the devices. For example, in Figure 5-3 only UUT, Device 1 and Device 2 need to be functioning devices; the other devices can be nonfunctioning representations of the UUT. The test fixture's lid is attached, and the test fixture is checked to ensure air flow is not leaking through its joints, cable egress locations, or at interface with the flow chamber.

The device under test and each adjacent identical device is operated to its rated maximum TDP level using an appropriate device stress test procedure. The device manufacturer should provide sufficient detail including software, firmware, and hardware required such that system integrators could recreate the same results if provided the same device. The following figures beginning with Figure 5-2 illustrate the test fixtures for E3 devices.

The current spec guarantees consistency in measurements but care should be exercised when using airflow impedance data for platform integration. Latches, backplanes, carriers, and cables will have a profound effect on both thermal and airflow impedance measurements and are not accounted for in this specification {and latches, backplanes, carriers and cables are not included in the AFI test methodology and only necessary cables are allowed in designated keep-in areas for MaxTherm and DTherm measurements} as they will vary from implementation to implementation. In addition, comparisons of airflow impedance across form factors should not be done as the test fixtures vary greatly.

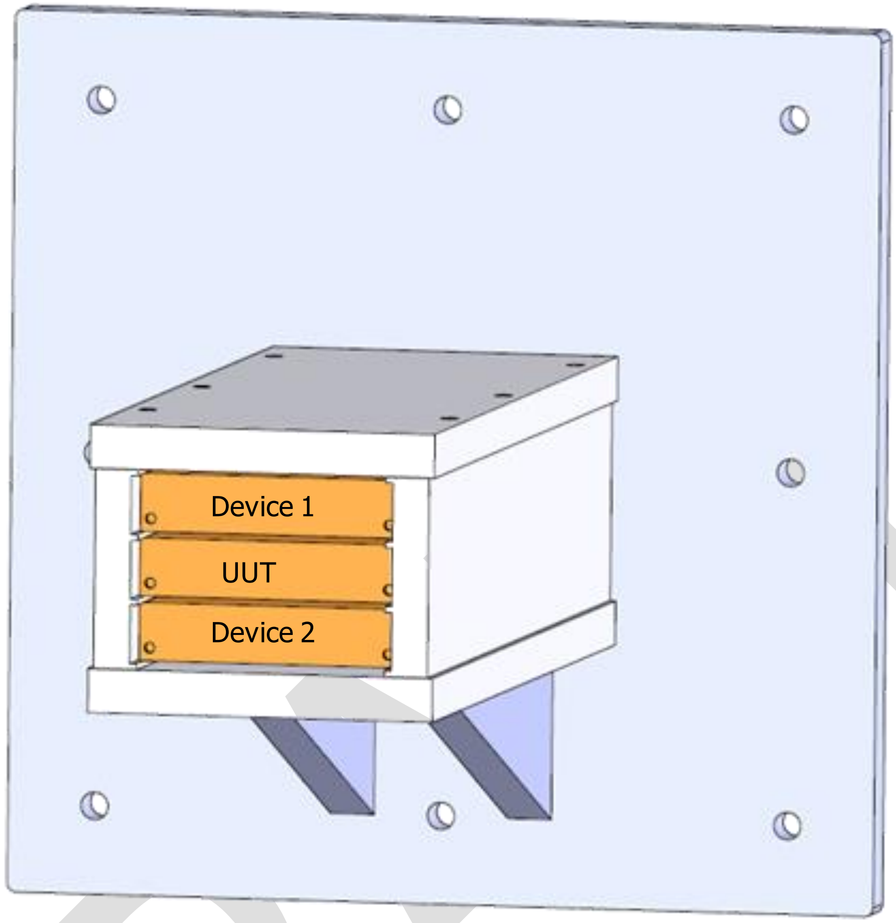

Figure 5-2: Example E3 2T Test Fixture and Device Setup

Developer's Note: *It is recommended that test fixtures be fabricated using materials with a thermal conductivity less than 1.0$Wm^{-1}K^{-1}$. This and the following examples show the devices in a horizontal orientation. The orientation of the devices may be horizontal or vertical.*

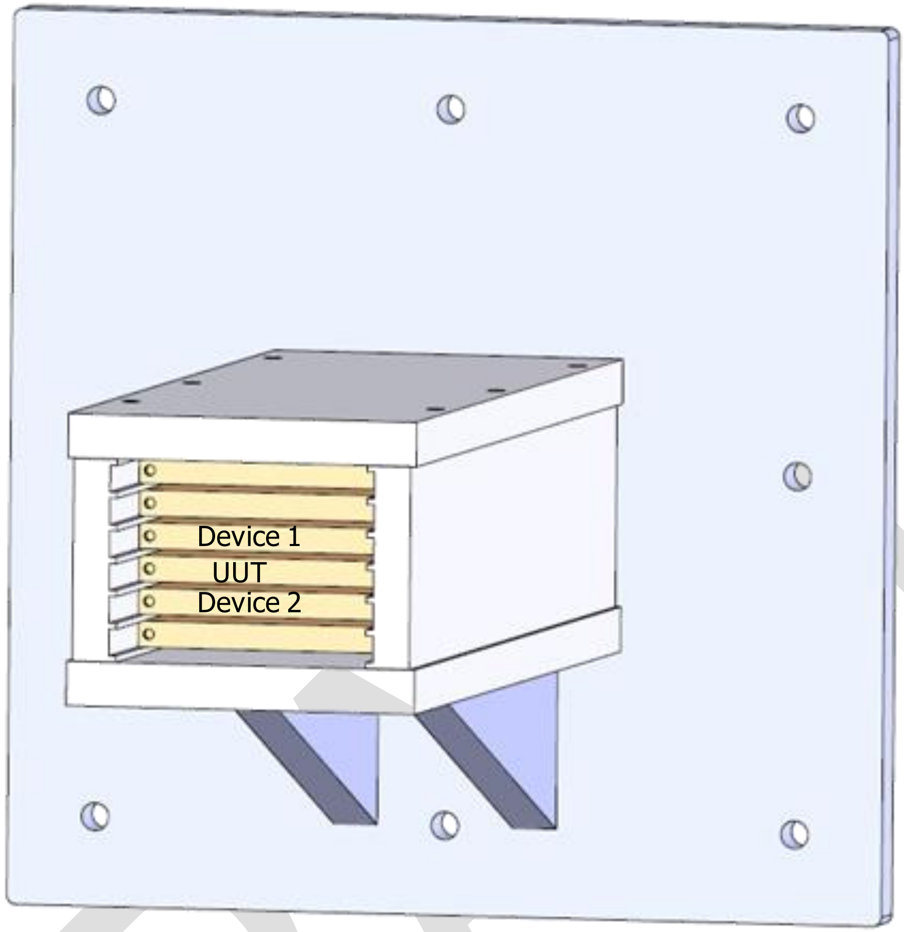

Figure 5-3: Example E3 1T Test Fixture

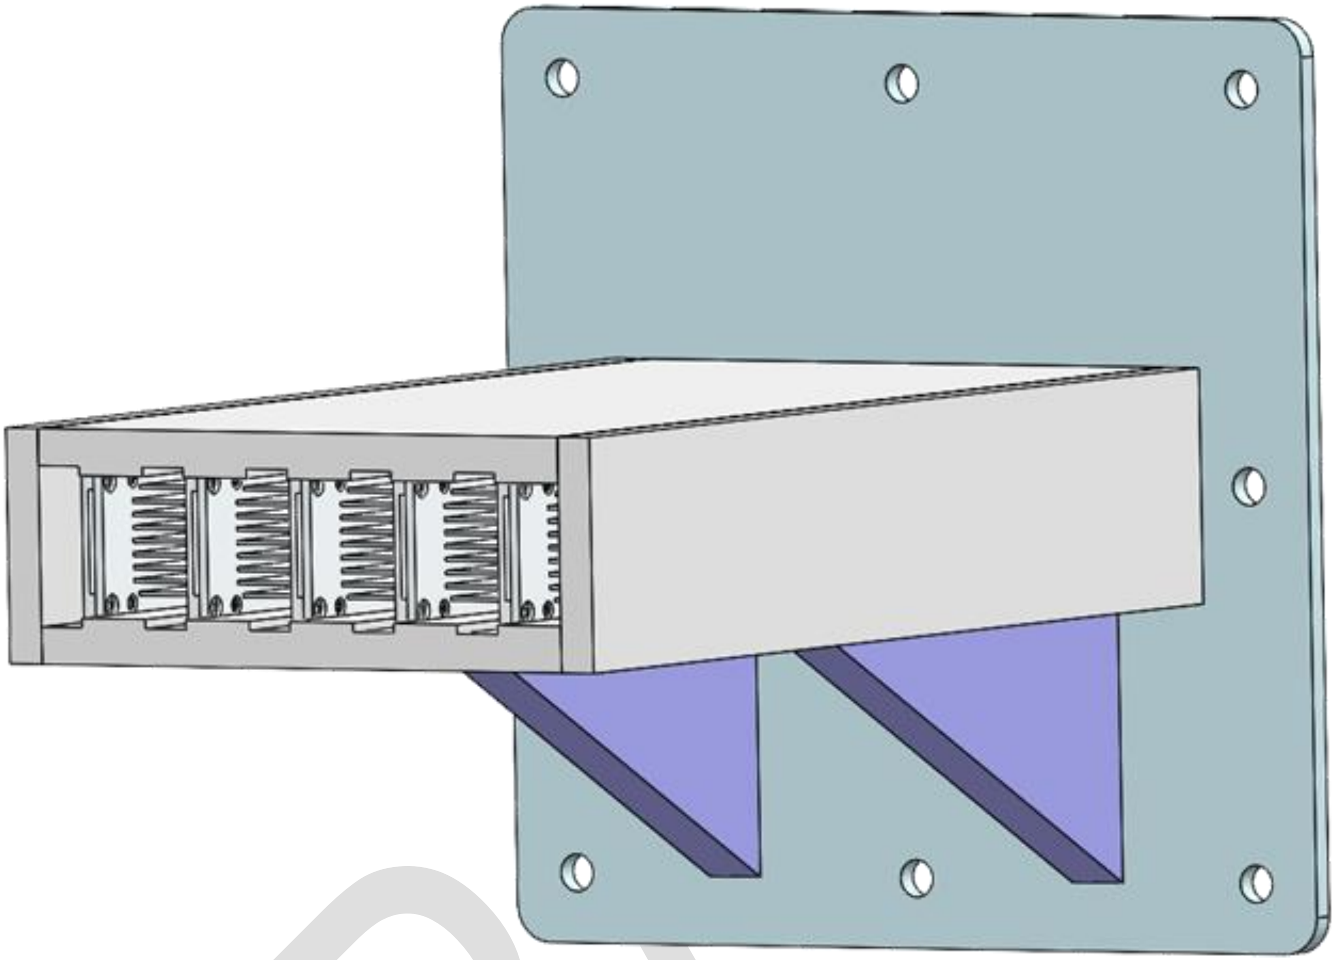

Figure 5-4: Example E1 Test Fixture

The test fixture design features are specified to ensure interoperability with E1 devices and consistent measurement. The following figures illustrate these interoperability points.

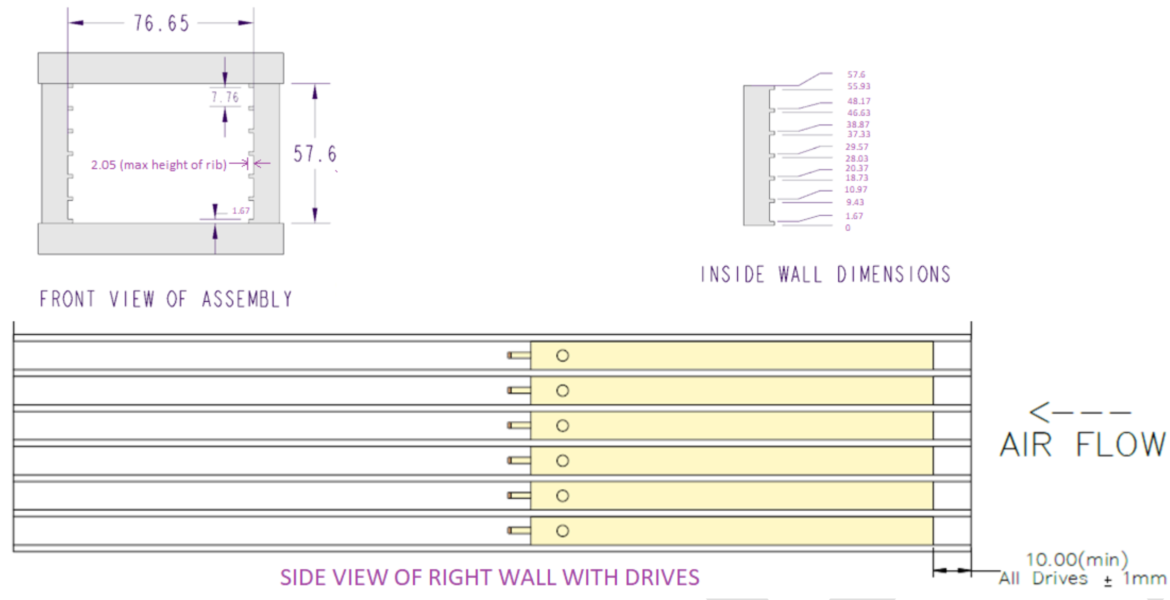

Figure 5-5: Required Dimensions of E3 1T Test Fixture

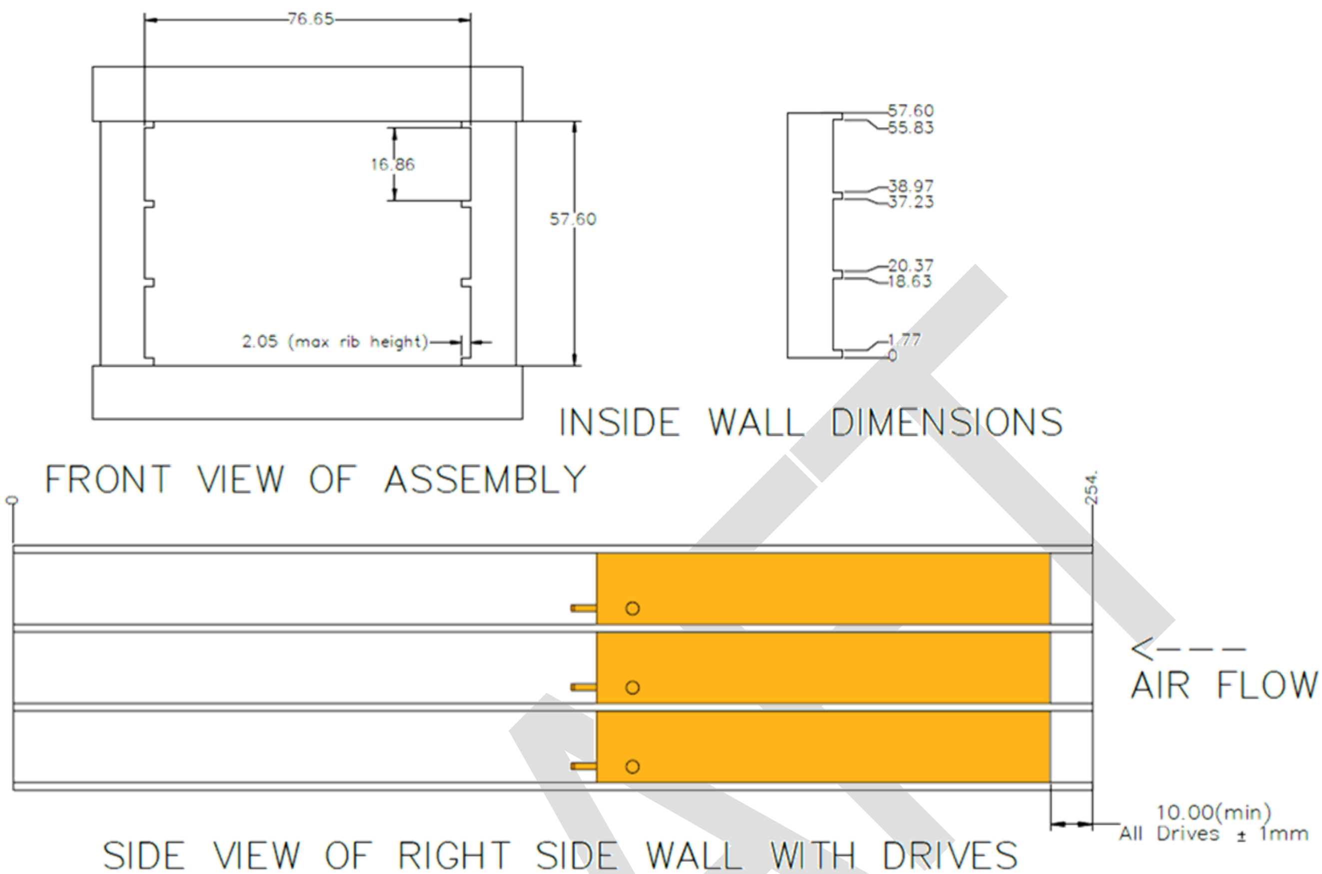

Figure 5-6: Required Dimensions of E3 2T Test Fixture

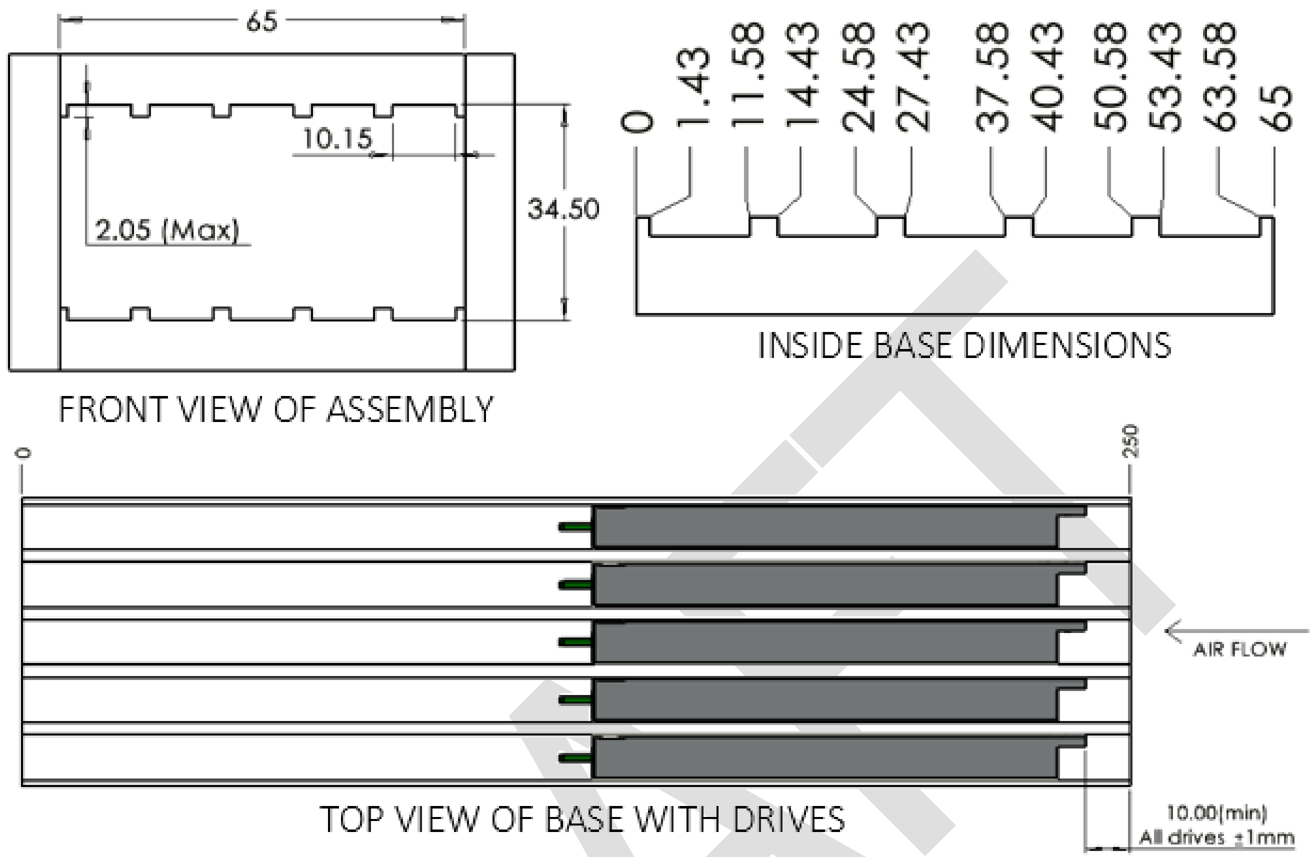

Figure 5-7: Required Dimensions of E1.S 9.5mm Test Fixture

Note: *E1 test fixtures end gap is ½ of the gap between devices, which differs from the E3 test fixture.*

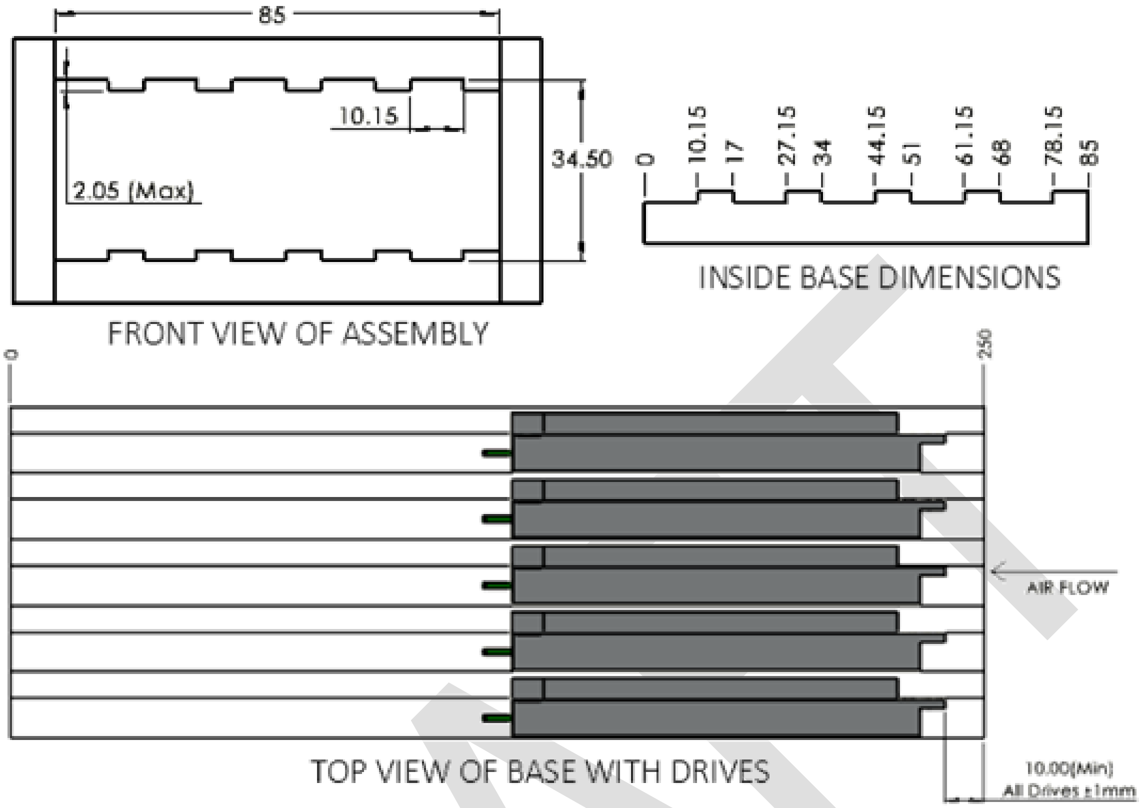

Figure 5-8: Required Dimensions of E1.S 15mm Test Fixture

Note: *The enclosure for this SSD is asymmetric about the PCB. The intent of this fixture is to prevent air from passing by the left most wall and SSD. This can be accomplished by using a gasket to seal the channel that would be between the left-most SSD and the fixture wall. This is done to simplify the CFM/SSD Channel calculation and promote balanced airflow impedance across the row of SSDs. E1 test fixtures end gap is ½ of the gap between devices, which differs from the E3 test fixture.*

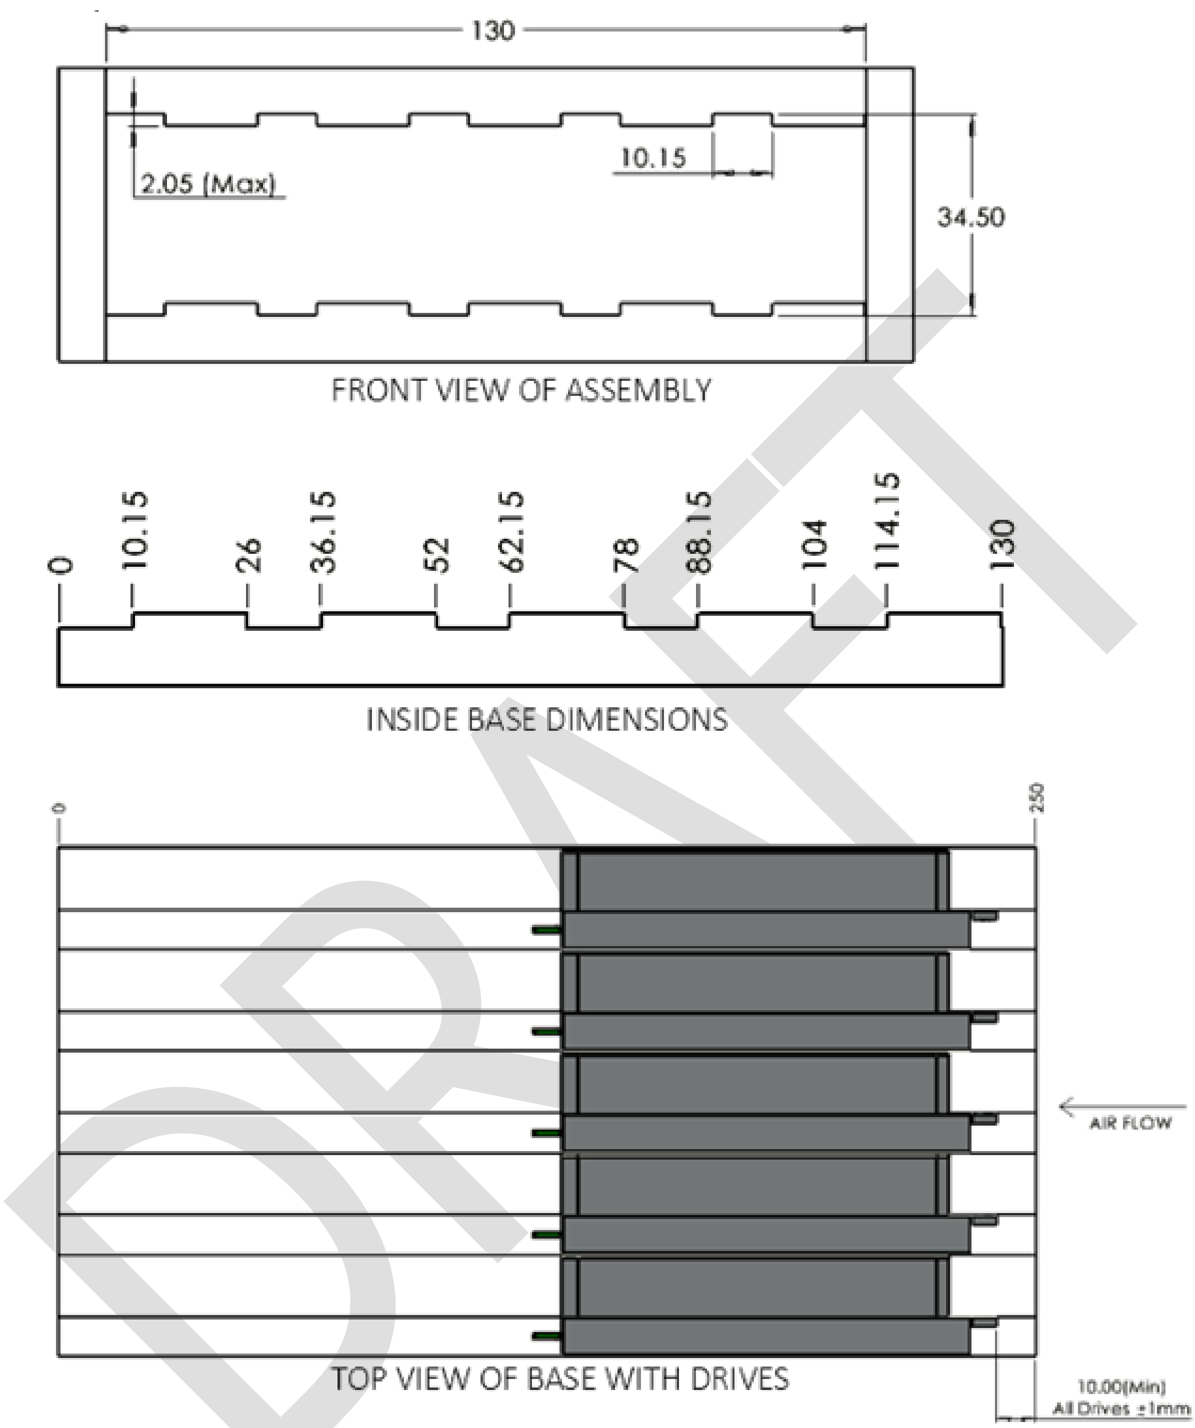

Figure 5-9: Required Dimensions of E1.S 25mm Test Fixture

Note: *The enclosure for this SSD is asymmetric about the PCB. The intent of this fixture is to prevent air from passing by the left most wall and SSD. This can be accomplished by using a gasket to seal the channel that would be between the left-most SSD and the fixture wall. This is done to simplify the CFM/SSD Channel calculation and promote balanced airflow impedance across the row of SSDs. E1 test fixtures end gap is ½ of the gap between devices, which differs from the E3 test fixture.*

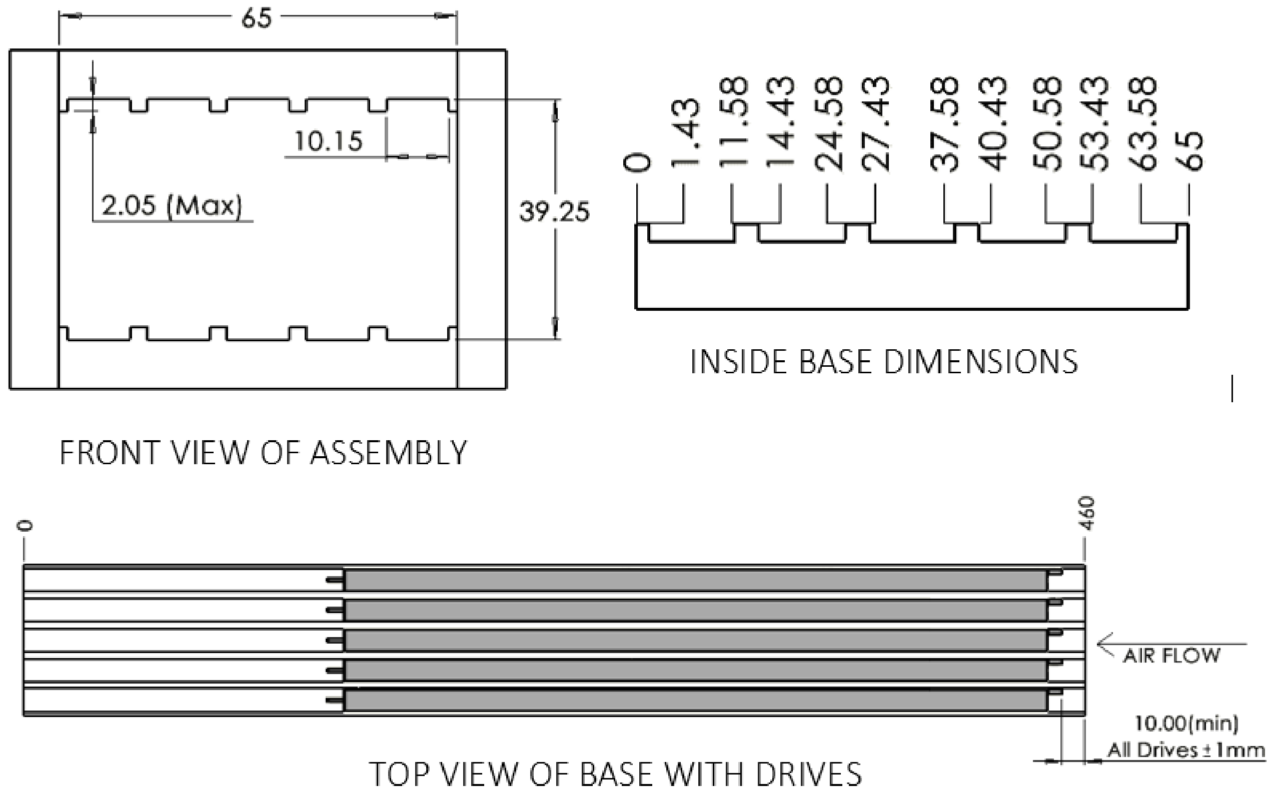

Figure 5-10: Required Dimensions of E1.L 9.5mm Test Fixture

Note: *E1 test fixtures end gap is ½ of the gap between devices, which differs from the E3 test fixture.*

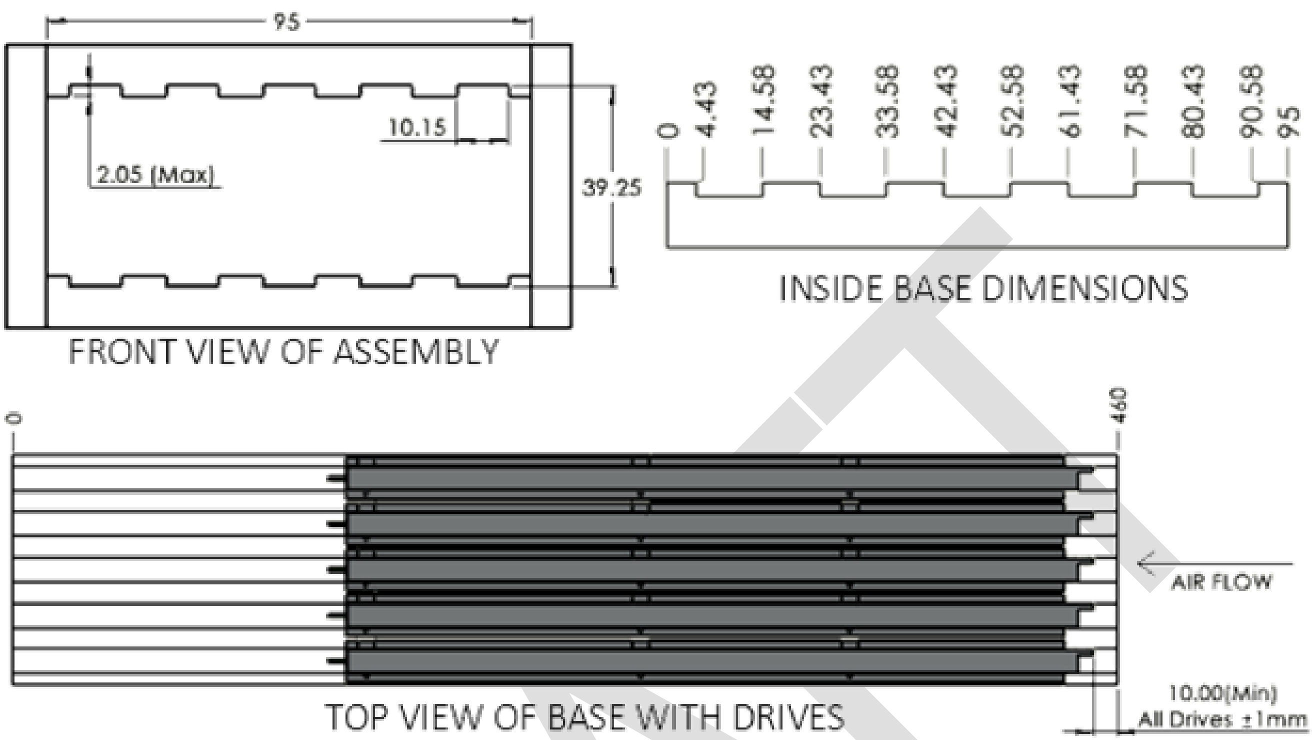

Figure 5-11: Required Dimensions of E1.L 18mm Test Fixture

A system board with enough slots to operate each device in the fixture at maximum link width and maximum link speed is needed for the testing. The device and system board shall be configured to operate at their maximum performance level. The platform power supply shall be capable of supplying the power required for the system board and the device to run at their highest performance levels. The following figures illustrate a cabling keep in for connector to devices.

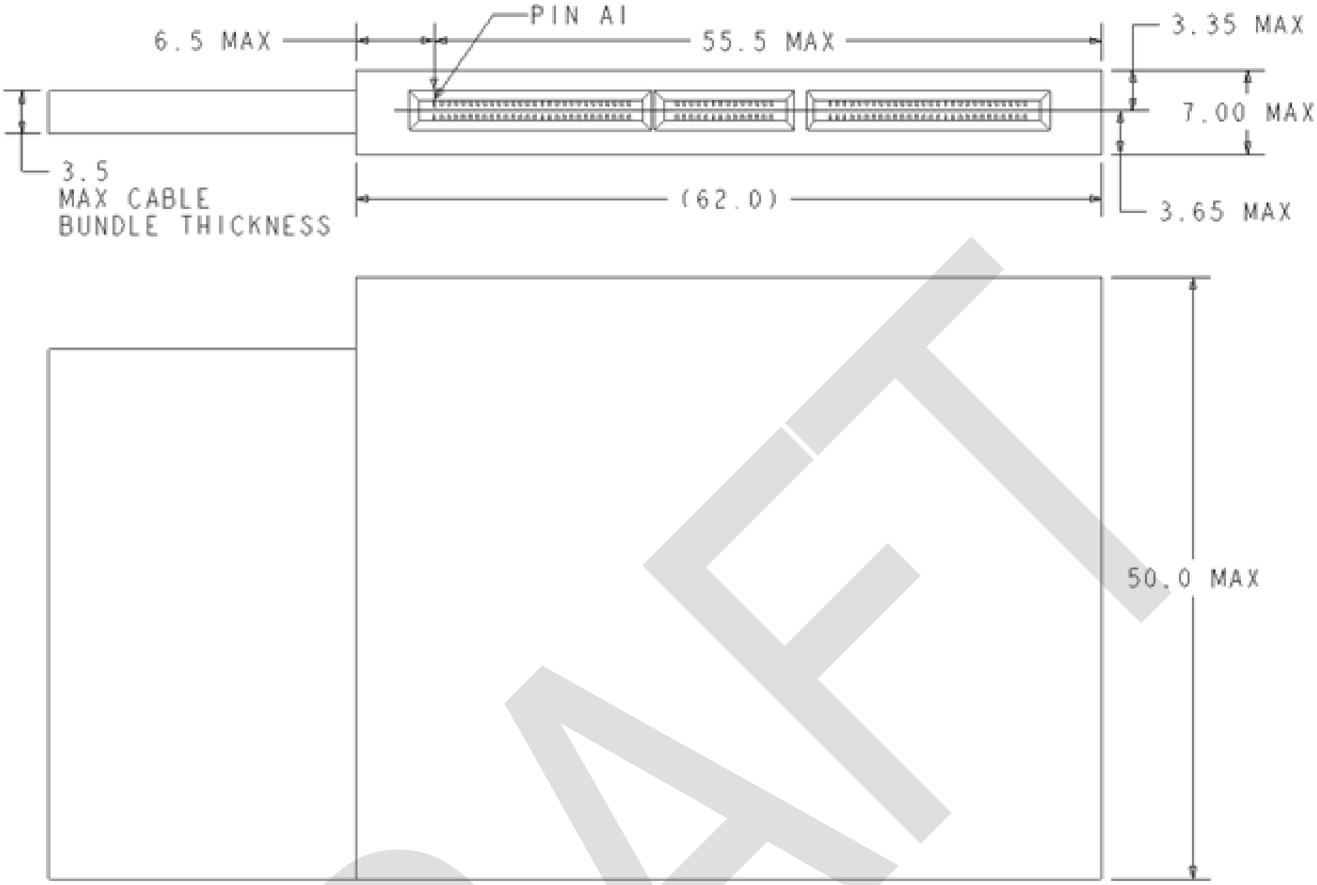

Figure 5-12: Maximum Test Cabling Keep-in Dimensions for E3

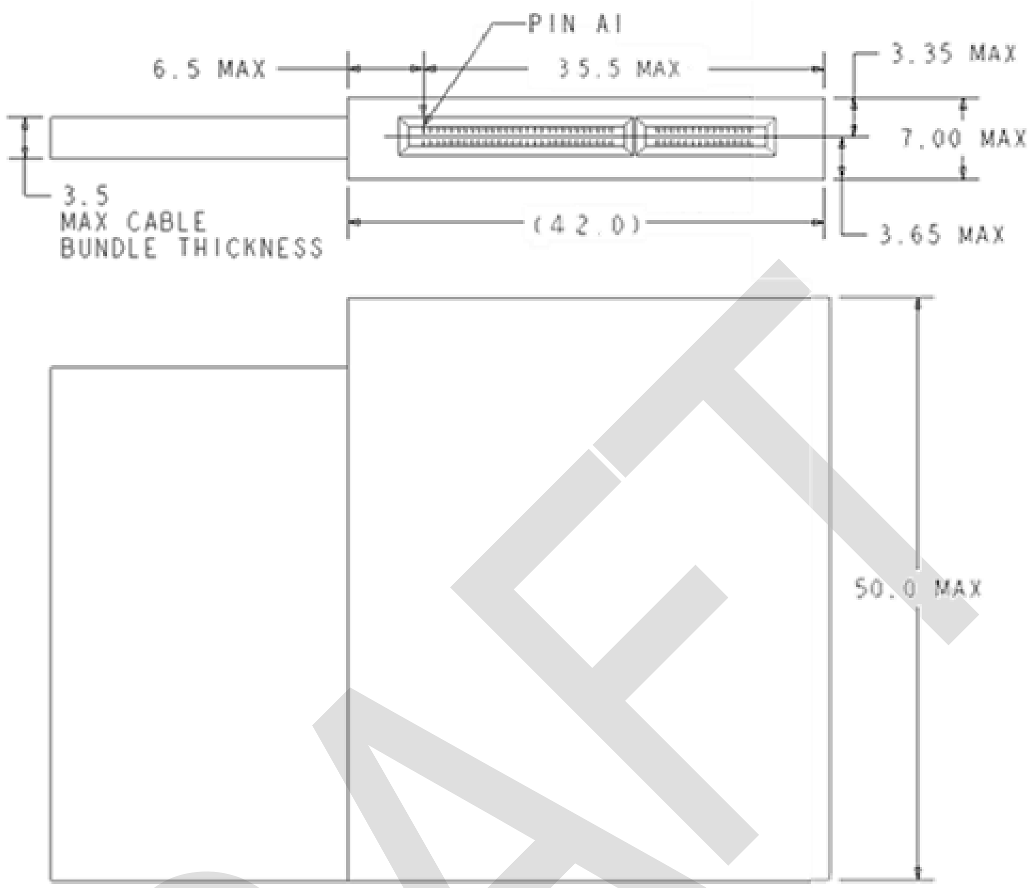

Figure 5-13: Maximum Test Cabling Keep-in Dimensions for E1

For MaxTherm and DTherm testing, the quantified volumetric airflow per device is reported.

At least one thermocouple sensor shall be placed to collect inlet air temperature data no less than 5mm from the inlet and no more than 4mm from the front of the device under test. Another thermocouple sensor shall be placed on the device under test to monitor critical local temperature. The location of local temperature measurement is defined by the implementer and should be reported in addition to thermal parameters. Figure 5-13 illustrates an example test setup connected to a flow bench.

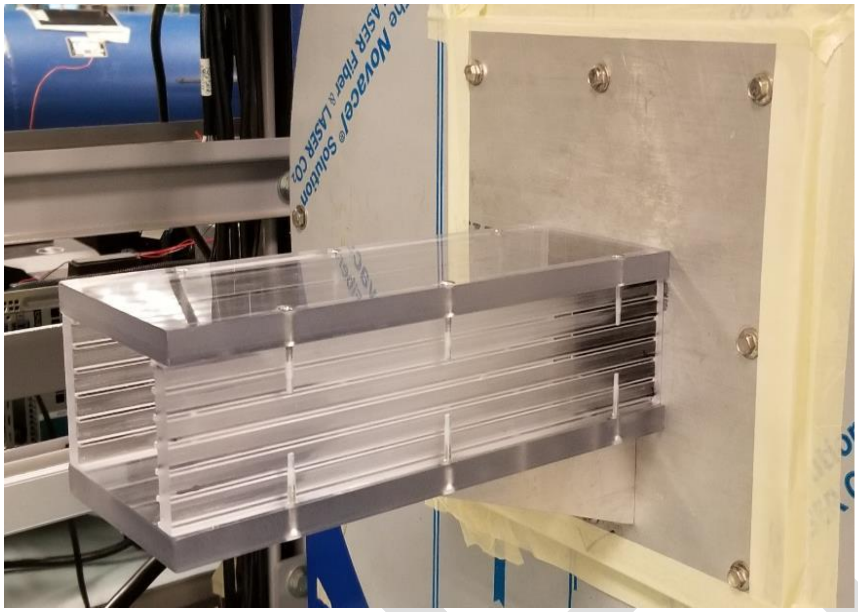

Figure 5-14: Actual Thermal Test Fixture Set-up

Note: SSD connectors downstream. Data cables not installed.

To evaluate the device cooling parameters (MaxTherm and DTherm), the approach air temperature can be scaled experimentally or analytically. Experimental air temperature scaling requires a heated flowbench, but analytical temperature scaling can be performed at room temperature. To perform analytical temperature scaling, use the equation as follows:

Equation 5-3: $$\boldsymbol{Scaled\ Device\ Temperature = (Desired\ Scaled\ Approach\ Temperature - Approach\ Air\ Temperature) + Device\ Temperature}$$

For Example, if the thermal limiting component on the UUT (Unit Under Test) is at 48 °C when flowrate is held constant at 2 CFM/dev, and the approach air temperature is 24 °C, we could scale the device temperature up as if the approach air temperature was 40 °C (or any other value chosen)

Example: $$\boldsymbol{64\ ^\circ C = (40\ ^\circ C - 24\ ^\circ C) + 48\ ^\circ C}$$

The device temperature at 2 CFM/dev and 40 °C approach would be 64 °C in this instance. Using this method, and the thermal temperature limit for the limiting component, the theoretical maximum approach temperature can be calculated at each flowrate for the different workloads associated with MaxTherm and DTherm.