Scaling Enterprise SSD Assembly for the AI Era: Why Single-Part Thermal Gels Win on Throughput

A Deep Materials White Paper on Low-Oil-Seepage Thermal Gels for EDSFF Enterprise SSDs

Rev 1.0

1. Introduction: AI Demand Is Turning eSSD Assembly Throughput Into a Bottleneck

The rapid buildout of AI infrastructure is driving an unprecedented surge in demand for enterprise SSDs (eSSDs). Training and inference clusters require massive, fast local and networked storage, and hyperscalers and OEMs are scaling eSSD production volumes accordingly. For manufacturers, this shift is not just about building more drives — it is about building them faster, without sacrificing the thermal reliability that high-power NAND and controller packages demand.

Thermal interface material (TIM) selection and application method are now a direct lever on assembly line throughput, not just a thermal-performance decision. This white paper argues that the choice between single-part and two-part thermal gels is the single most consequential TIM decision an eSSD manufacturer can make for production rate, and lays out a practical path to higher throughput:

- **Eliminate two-part gels where possible:** removing the mixing, working-life, cure-station, energy-use, and rework constraints that cap line speed and complicate servicing.
- **Move to single-part gels dispensed faster** — single-part gels are ready to use as dispensed, with no cure step gating assembly.
- **Scale further with screen printing** — applying gel to dozens of locations in a single shot, well beyond what point-by-point dispensing can achieve.

Gap pads remain in use in some designs, but as discussed below they are largely a secondary option for high-density, high-throughput eSSD lines because they resist automation. The remainder of this paper focuses on the single-part vs. two-part gel decision, the throughput and reliability trade-offs each path creates, and how Deep Materials' low-oil-seepage single-part gels remove the historical objection to running single-part chemistry at scale.

2. EDSFF: The Form Factor Driving the Volume

The EDSFF (Enterprise and Datacenter Standard Form Factor) is purpose-built for enterprise SSDs, optimizing storage density, thermal efficiency, and performance in data center environments. Compared to legacy M.2 and U.2 form factors, EDSFF offers superior scalability and thermal headroom for high-performance storage. Three main variants exist: E1.S (targeted at 1U servers), E1.L (optimized for 1U storage servers with higher byte density for cloud infrastructure), and E3 (in four sub-variants — Long Thick, Short Thick, Long Thin, and Short Thin — designed for 2U servers and storage). All EDSFF variants share the SFF-TA-1002 connector and use the PCIe or NVMe protocol.

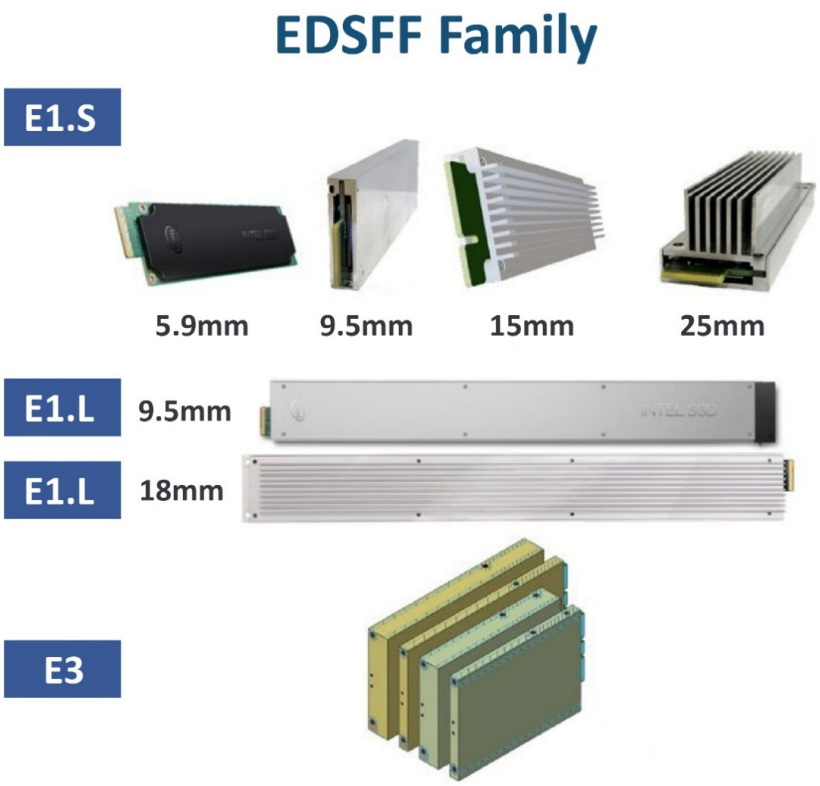


Figure 1: EDSFF form factors for enterprise SSDs (Source: allion.com)

EDSFF's design point — higher power density (up to 40W or more per drive), larger surface area for airflow, and support for hot-swapping — is exactly what makes thermal management, and by extension TIM selection, a first-order design and manufacturing concern rather than an afterthought.

2.1 Why Thermal Management Cannot Be an Afterthought

Effective thermal management prevents throttling, reliability loss, shortened SSD life, and data-integrity risk.

Enterprise SSD enclosures are typically aluminum or magnesium alloy, chosen for structural integrity and thermal conduction. The internal layout is optimized for airflow, with thermal interface materials bridging heat-generating components to the enclosure, which functions as the heat sink.

3. TIM Options: Gels as the Primary Path, Pads as a Secondary Option

Gap fillers are needed at multiple locations on an eSSD — NAND devices, controllers, power converter ICs — to bridge components to the heat sink. Historically, gap pads served this role. For high-volume eSSD manufacturing, however, gels are the fundamentally better solution, and pads are best treated as a secondary or fallback option rather than the default.

The core limitation of gap pads is that they resist automation: pads are cut, picked, and placed as discrete parts, which is inherently harder to automate at speed than dispensing a flowable material. In practice this means pad-based assembly lines are more dependent on manual or semi-manual placement, capping throughput exactly where AI-driven volume growth demands the opposite.

- **Automation:** Gels can be dispensed or screen-printed by automated equipment, directly improving production efficiency and consistency at high volume — pads cannot be applied this way.
- **Conformability:** Gels are highly conformable and can be dispensed in various shapes and patterns, filling irregular gaps and improving surface contact; pads are fixed-geometry parts.
- **Reliability:** Compressed pads transmit stress, shock, and vibration from the enclosure to the die, risking damage to delicate packages. Pads also tend to stiffen as their conductivity increases, compounding this risk.
- **Thermal Performance:** Gels generally provide superior thermal conductivity and lower thermal resistance across uneven or textured surfaces.
- **Cost at Scale:** Dispensing equipment carries a higher upfront cost, but gels become more cost-effective than pads in large-scale production due to reduced labor and faster application times.

Given these advantages, the real decision eSSD manufacturers face is not “gel or pad” — it is which type of gel chemistry to run, since that choice determines how much of the automation and throughput upside above is actually realized on the line.

4. The Core Decision: Single-Part vs. Two-Part Thermal Gels

Thermal gels come in two fundamental chemistries, and this distinction — more than any other TIM choice — determines achievable line throughput.

- **Single-part gels are ready-to-use, non-curing materials** dispensed in their final usable state. They are highly viscous and stable, meaning assembly can happen immediately after dispensing or long afterward, with no chemical change in between.
- **Two-part gels** cure via a chemical reaction (temperature or humidity driven) after being dispensed on the application surface. They are not in their final state at the point of dispensing — they need time, and often heat, to become stable.

4.1 Two-Part Gels Impose a Structural Throughput Ceiling

Two-part chemistry is common in the industry today, but it comes with manufacturing costs that directly cap production rate:

- **Limited assembly window:** Parts must be assembled within a narrow working-life window after dispensing, before the gel begins to set. This forces a rigid, tightly sequenced process flow.
- **A mandatory curing station and added energy use:** Full cure typically requires heating above ambient temperature, which means dedicating physical staging or oven space on the line. The oven adds cycle time, floor space, capital equipment, and ongoing energy consumption - all costs that do not exist for single-part chemistry.
- **Room-temperature cure does not eliminate the constraint:** Although some two-part gels can cure at ambient conditions, the process can take on the order of 24 hours and may be less uniform because it is sensitive to local temperature, humidity, and bond-line geometry. During this hold period, partially cured material can remain mobile, increasing the risk of oil migration or leakage. The product also cannot be safely handled or shipped as a finished assembly while material may still move within the enclosure.
- **Difficult rework after cure:** Once a two-part gel crosslinks, it can adhere to both mating sides of the assembly and form a cohesive bridge between them. Opening the SSD enclosure for inspection, repair, or component replacement can then require substantial separation force and cleanup, with added risk of damaging the board, components, or casing.
- **Cure-time variability:** Small lot-to-lot variation in the curing catalyst from the supplier can shift cure time, making line balancing and takt time harder to hold constant.
- **Shelf-life constraints:** Two-part gels typically carry a shelf life of only about six months, adding inventory and planning overhead.

Crucially, the curing step is not just a throughput cost — it is also a reliability risk. If curing is incomplete or inconsistent (due to under-heating, catalyst variation, or a rushed process flow used to protect line speed), two-part gels can end up under-cured. An under-cured two-part gel

behaves in some respects like a poor single-part gel: it can still bleed oil, and unreacted material can migrate or remain mechanically unstable, undermining the very stability advantage two-part chemistry is chosen for in the first place. In other words, cutting corners on cure time to protect throughput can quietly reintroduce the oil-bleed problem discussed in Section 5 while also sacrificing dimensional stability.

4.2 Single-Part Gels Remove the Cure Step Entirely

Because single-part gels are already in their final usable state at the point of dispensing, they eliminate the throughput constraints above outright:

- **Faster dispensing, no waiting:** There is no working-life clock running after dispensing, so line speed is limited by the dispensing equipment itself, not by chemistry.
- **Lower-energy process, with no curing station required:** Removing the cure step eliminates oven power, staging space, added cycle time, and associated capital equipment. This creates a lower-energy manufacturing path and can reduce the process-related carbon footprint of each assembled drive.
- **No ambient-cure hold or shipping delay:** Because the gel is already in its usable state, the assembly does not need to remain stationary for a day while cure progresses. It can move directly to the next operation or shipment without the risk of partially cured material shifting inside the enclosure.
- **Easier rework and servicing:** A non-curing single-part gel remains compliant rather than forming a crosslinked bridge between mating surfaces. Enclosure components can be separated more readily for inspection or repair, and the material can be cleaned and reapplied without the force and cleanup associated with breaking apart a cured two-part gel.
- **Assembly flexibility:** Single-part gels can be assembled immediately after dispensing or much later, decoupling the dispensing step from final assembly and simplifying scheduling.
- **Pre-application to the heat sink:** Because there is no cure clock, single-part gel can be applied to the heat sink or case in advance — including via screen printing prior to shipment — rather than only at the point of final SSD assembly.

This is the throughput path outlined in Section 1: removing the two-part cure bottleneck, dispensing single-part gel faster, and ultimately moving to screen printing (Section 6) to apply gel to many locations in a single step. Historically, the objection to running single-part gel at high volume has been oil bleed — which is the problem Deep Materials' low-oil chemistry is built to solve.

5. The Historical Objection: Oil Bleed in Single-Part Silicone Gels

Almost all silicone-based single-part thermal gels leach a measurable amount of silicone oil over time. Silicone oil has one of the lowest surface tensions of any common fluid, which is precisely what gives it excellent spreadability and dispensability — but that same property drives it to seep out of the gel matrix and spread across surrounding electronics. This oil migration, or "bleed," degrades thermal performance over time and is a well-known reliability issue for silicone-based TIMs generally.

As noted in Section 4.1, oil bleed is not exclusive to single-part gels: an under-cured two-part gel can exhibit similar migration and stability problems, since incomplete curing leaves unreacted, mobile material behind. Addressing oil bleed at the chemistry level therefore benefits both the single-part throughput path this paper recommends and any two-part processes still in use.

- **Contamination of electronics:** Leaked oil can spread onto nearby components and connectors, degrading electrical connections and increasing the risk of short circuits or elevated resistance. This is a particular concern near optical sensors or cameras; in extreme cases, liquid droplets have been observed forming at the periphery of the SSD.
- **Reduced thermal performance:** As a gel loses oil, it can dry out and crack, degrading contact between the component and heat sink, raising thermal resistance and cutting cooling efficiency.
- **Aesthetic and mechanical issues:** Oil residue can mar the appearance of high-end systems and attract dust and contaminants, adding to customer complaints.



Figure 2: Silicone oil migration observed in a real product over time (shown here from a pad; the same migration mechanism applies to gels).

This oil-bleed risk is the reason some manufacturers have defaulted to two-part gels despite their throughput costs — accepting a slower, cure-gated process in exchange for perceived stability. As shown next, this trade-off is no longer necessary.

6. Removing the Trade-Off: Deep Materials Low-Oil Single-Part Gels

Using a proprietary, patented formulation method, Deep Materials has developed a single-part gel family that demonstrates ten times lower oil migration than conventional single-part gels — removing the historical reason to fall back on slower two-part chemistry. The product family includes TCG-Gel 4000 LO and TCG-Gel 6000 LO, with thermal conductivities as high as 6 W/m·K.

Oil migration is measured by heating a known gel quantity on absorbent grid paper and tracking the resulting oil stain over time.

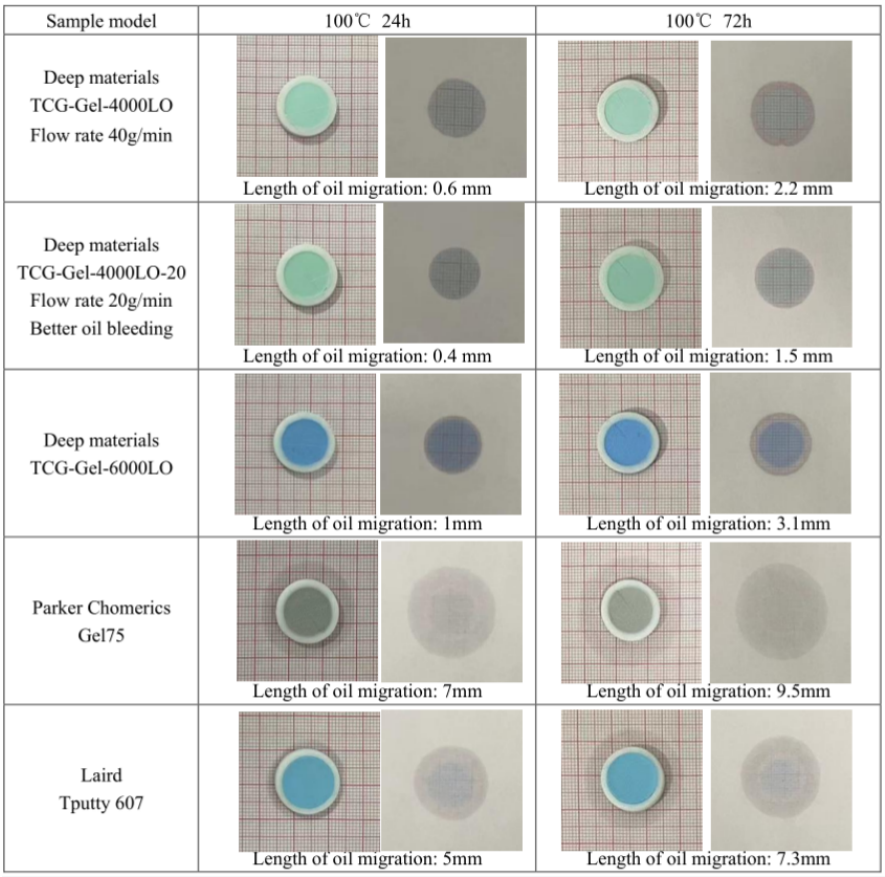


Figure 3: Oil seepage comparison — Deep Materials low-oil formulation vs. industry-standard competitor products.

Table 1 summarizes typical properties. Deep Materials low-oil gels are available in multiple variants, including 4 and 6 W/m·K thermal conductivities; the 4 W/m·K gel is also available in a very-low-volatility option. Additional variants are under development.

Property	Value	Test Method
Color	Green	Visual
Flow Rate (g/min)	40	30 cc tube, 0.090" nozzle, 90 psi
Thermal Conductivity (W/m·K)	4.0	ASTM D5470
Minimum BLT (µm)	70	—
Density (g/cc)	3.2	ASTM D792
Breakdown Voltage (kV/mm)	>8.0	ASTM D149
Operating Temperature Range (°C)	-50 to 200	—
Volatility (%)	0.15	ASTM E595
Flammability Rating	V-0	UL94
Volume Resistivity (Ω-cm)	>10^13	ASTM D257
Heat Capacity (J/g·K)	1.1	ASTM E1269
Oil Bleeding Size (mm)	< 1 mm	3-gram sample, 14 mm diameter, 100°C for 24 hours on grid paper

Table 1: Properties of Deep Materials TCG-Gel 4000 LO



Figure 4: Dispensed single-part low-oil gel inside an SSD.

Several enterprise SSD manufacturers have integrated thermal gels into their products. Most notably, Deep Materials has commercialized this low-oil, single-part gel specifically for high-density applications such as modern enterprise SSDs — and even under close component proximity and sustained high-temperature use, no oil leaks.

7. Scaling Throughput: From Faster Dispensing to Screen Printing

With the oil-bleed objection resolved, single-part gel opens up a clear throughput escalation path: faster point dispensing first, then screen printing for a step-change in production rate.

7.1 Faster Dispensing With Single-Part Gel

Because single-part gel requires no post-dispense cure, automated dispensing equipment can run at the speed the hardware allows, without pausing for chemistry. Precision nozzles apply gel directly to designated board locations; equipment can be programmed for different patterns and volumes to match component layout and thermal needs, with high-speed vision systems detecting and correcting misalignment in real time.

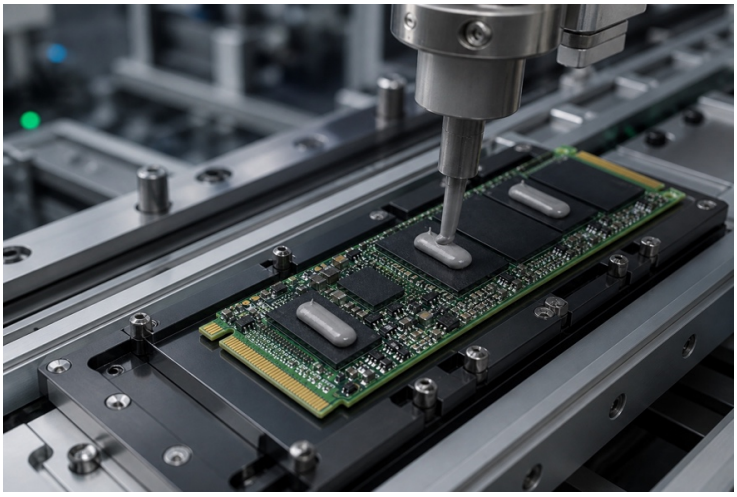


Figure 5: Automated dispensing of single-part gel.

Larger reservoirs (a 3 L pail vs. a 1 L syringe, for example) reduce changeover frequency, though larger pails require more pressure to move the highly viscous gel. Automated dispensing also improves consistency, reduces manual labor and material waste, and enables real-time quality control on gel thickness and placement — all without the cure-related constraints two-part chemistry would impose.

Even with these gains, as component density increases and the number of dispense locations per board grows, point-by-point dispensing throughput can still become the limiting step. That is where screen printing extends the throughput curve further.

7.2 Screen Printing: Applying Gel to 30+ Locations in One Shot

Screen printing is a well-established technique for applying thermal grease to heat sinks, and has been extended to gels in computing applications such as laptops, though it has seen limited use for SSD heat sinks to date. Deep Materials has extended this technique to apply gel to more than 30 locations on one side of a heat sink in a single shot, and has developed a semi-automatic screen-printing method for this purpose.

This approach requires the gel to be applied to the heat sink prior to shipping — which single-part chemistry makes practical, since there is no cure clock to manage in transit — rather than dispensed at the point of SSD assembly. As with pre-applied greases or gels shipped today on conventional PC or server heat sinks, standard handling precautions prevent smudging in transit. For manufacturers looking to move beyond dispensing-limited throughput, screen printing represents the next step-change in production rate. Contact Deep Materials for implementation details.

7.3 Assembly Steps (Either Application Method)

Assembly requires (1) identifying heat sources, (2) preparing surfaces, (3) applying gel by dispense or print, (4) assembling without displacement, and (5) verifying thermal performance.

8. Screen-Printed Gel on an SSD Case: Example

The sequence below shows a representative screen-printing process for applying thermal interface material through a stencil. Deep Materials adapts this process for SSD heat spreaders and cases: a stencil is aligned to the case (Figure 6), gel is applied across the stencil to achieve variable, location-specific coverage (Figure 7), the stencil is removed to leave the gel pattern in place (Figure 8), and the completed print is ready for assembly (Figure 9).

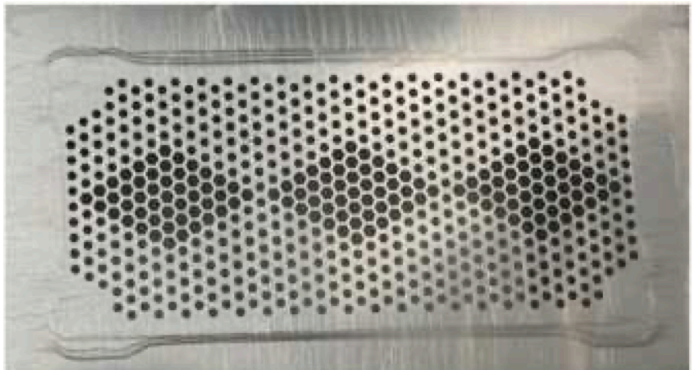


Figure 6: Screen-print stencil (similar to SSD case geometry).



Figure 7: Variable-coverage screen print application.

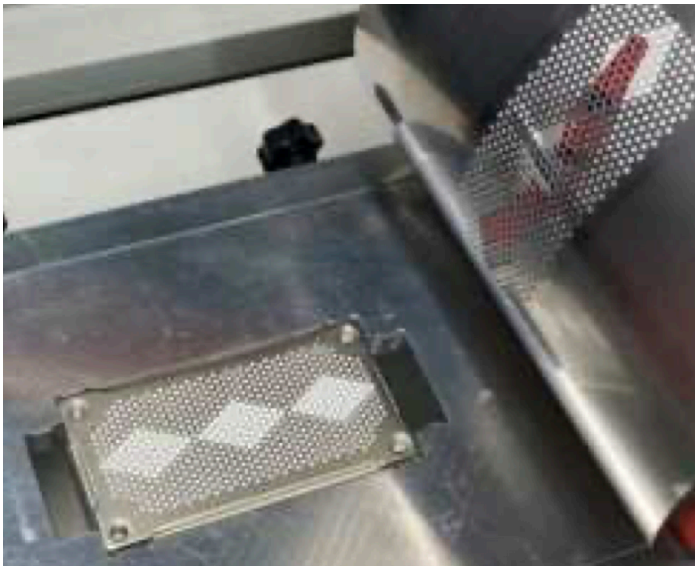


Figure 8: Screen print removal.

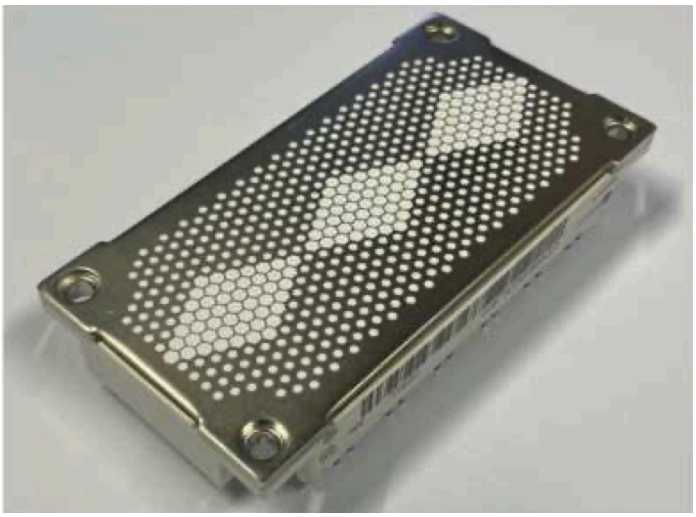


Figure 9: Completed screen printing (analogous to SSD case).

9. Conclusion

AI-driven demand for enterprise SSDs is pushing manufacturers to increase assembly throughput without compromising thermal reliability. Gap pads are a secondary option at best for this environment because they resist automation. Among gels, the single-part vs. two-part decision is the pivotal throughput lever: two-part chemistry imposes a structural ceiling through limited assembly windows, mandatory curing infrastructure, added energy use, long ambient-cure holds, difficult rework, cure-time variability, and short shelf life — and when curing is incomplete, it can reintroduce oil-bleed and stability problems of its own. Single-part gel removes the cure step entirely, enabling faster dispensing today and a clear path to screen printing for a further step-change in throughput. Deep Materials' patented low-oil single-part chemistry — demonstrating ten times lower oil migration than conventional single-part gels — removes the last major objection to running single-part gel at scale, making it the preferred path for high-throughput enterprise SSD manufacturing.

For more information about Deep Materials low-oil single-part materials or to request samples, please contact: sales@deep-materials.com

Steve Frayne Jr. — Steve.Frayne@Deep-Materials.com

Himanshu Pokharna, PhD — Pokharna@Deep-Materials.com

Figure 1 source: https://www.allion.com/server-series_form-factor-connector_02/

Note: Limited acrylic-based single-part gel options exist, but their longevity and weatherability are not as strong as silicone-based products, so uptake in the high-lifetime enterprise market remains limited and is not recommended.