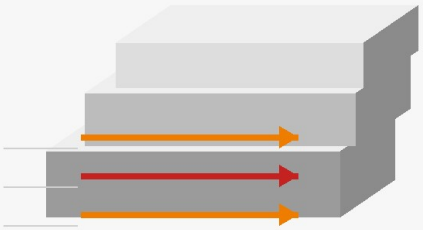

# Heat Distribution in Multilayer Clad Materials under Steady-State and Transient Loading

Technical whitepaper with statistics, renderings, and design logic


## Author

Revolit GmbH
Samstagernstrasse 41
8832 Wollerau
Switzerland



Edition April 2026

# Imprint and document status

Document framework, edition logic, and intended use

Document title: Heat Distribution in Multilayer Clad Materials under Steady-State and Transient Loading. Document type: technical whitepaper with statistics, renderings, and design logic. Author and publisher: Revolit GmbH, Samstagernstrasse 41, 8832 Wollerau, Switzerland. This edition is written as a strategic and technical working document for engineering, sales, quality, and management.

The whitepaper focuses on the physical core of the subject: resistance paths through the layer stack, boundary-side heat release, functional interlayers, and the time scales of thermal penetration. Statistical statements in Chapter 6 are based on a reduced Monte Carlo study with 12,000 virtual stacks. They are meant as engineering evidence for prioritization and robustness review, not as a product-specific warranty statement.

**Document status**

Edition April 2026 - internal whitepaper format for technical positioning, early sizing, and release governance.

# Legal and technical notice

Scope, limits of applicability, and release logic

This whitepaper supports technical classification and economically sound design decisions around heat distribution in multilayer clad systems. It does not replace project-specific sizing, material qualification, customer-specific approval, or normative acceptance. Whenever the application is safety-critical, certification-relevant, contract-defined, or highly process-sensitive, the actual boundary conditions, batches, bond states, and manufacturing route must be verified separately.

All values, equations, statistics, and worked examples are methodological aids. Thermal decisions are only reliable when load collective, geometry, bond quality, cooling connection, and manufacturing scatter are evaluated together. Interfaces and transiently loaded components in particular require a project-specific release chain. Recommendation: read this whitepaper together with the load definition, the material data sheet, and the test plan.

**Liability logic**

The nominal material label never substitutes for evidence on geometry, boundary condition, and bond continuity.

# Management Summary

Core technical statements and economic levers

In multilayer clad systems, heat distribution is not determined by a single material property. The result follows from the complete path: series resistance through thickness, lateral spreading in conductive top layers, local temperature jumps at bond lines, and the quality of coupling to the sink side. For management this leads to a direct priority sequence: first classify the load case and the dominant heat path, then scale the required model and test depth accordingly.

The commercial value is substantial. Early identification of interface-sensitive designs shortens development loops, reduces wrong turns in material and process selection, and makes thermal risk visible during quotation and concept review. In scalable products, robustness against scatter matters more than a strong nominal point. Good thermal architecture therefore improves both technical performance and release reliability.

- Interface quality is often the fastest lever on hot-spot risk and robustness.
- A highly conductive cap layer smooths the field, but it never compensates for a poor bond.
- Transient duty must be judged with α_eff, penetration depth, and t95 - not with k alone.
- Production-capable stacks differ less in the nominal point than in their scatter resilience.

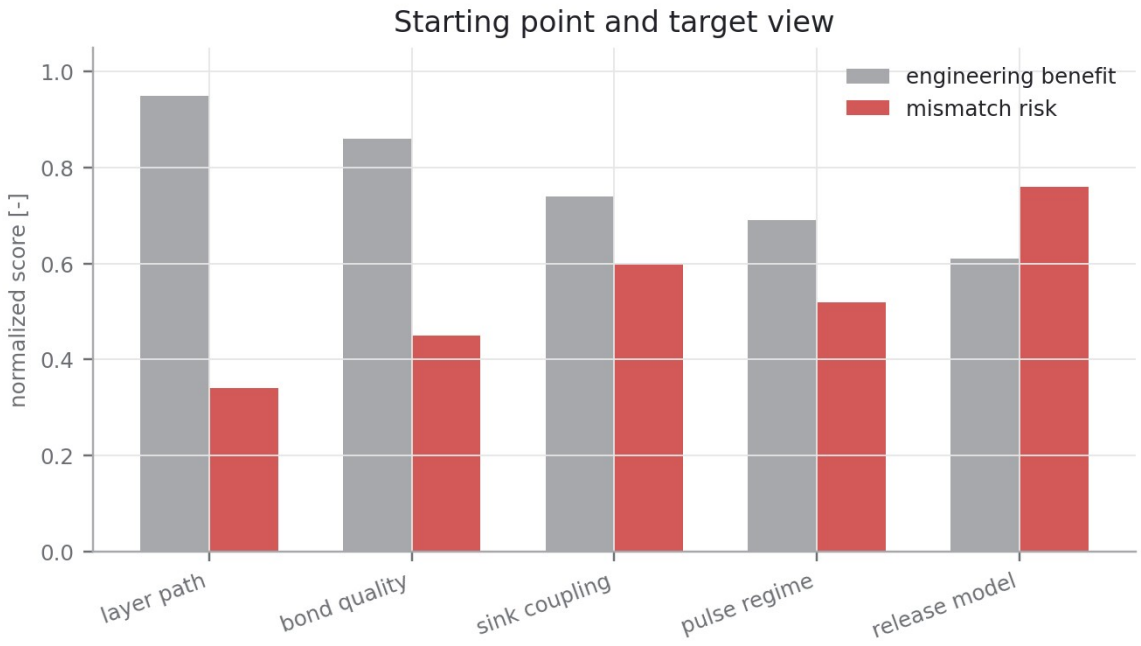


Figure 1: Condensed management view of benefit, mismatch risk, and review priority in the sizing space.

# Executive Summary - decision system

Four levers that govern project success

This whitepaper compresses heat distribution in clad materials into an operating decision system. The key is to derive governance from four interacting layers: heat-path architecture, bond state, boundary-side release, and time scale. Teams that treat these dimensions separately will usually underestimate hot-spot risk, measurement effort, and the amount of evidence required for release.

The first lever is the interface: thin functional layers can dominate the total temperature rise. The second lever is the process route and its scatter: production capability is itself a thermal property. The third lever is the coupling of steady-state and transient assessment: the same stack behaves fundamentally differently under pulses than under continuous duty. The fourth lever is model economy: a digital twin is only economical when hand calculation, reduced model, field model, and measurement are connected cleanly.

- Classify the load and the dominant path before debating materials.
- Treat bond quality and sink coupling as release-critical variables.
- Escalate from screening to 2D/3D only when the regime truly demands it.
- Tie procurement language to measurable thermal criteria rather than to labels only.

### Thermal decision matrix

Model and test depth should scale with load severity, path complexity, and release risk.

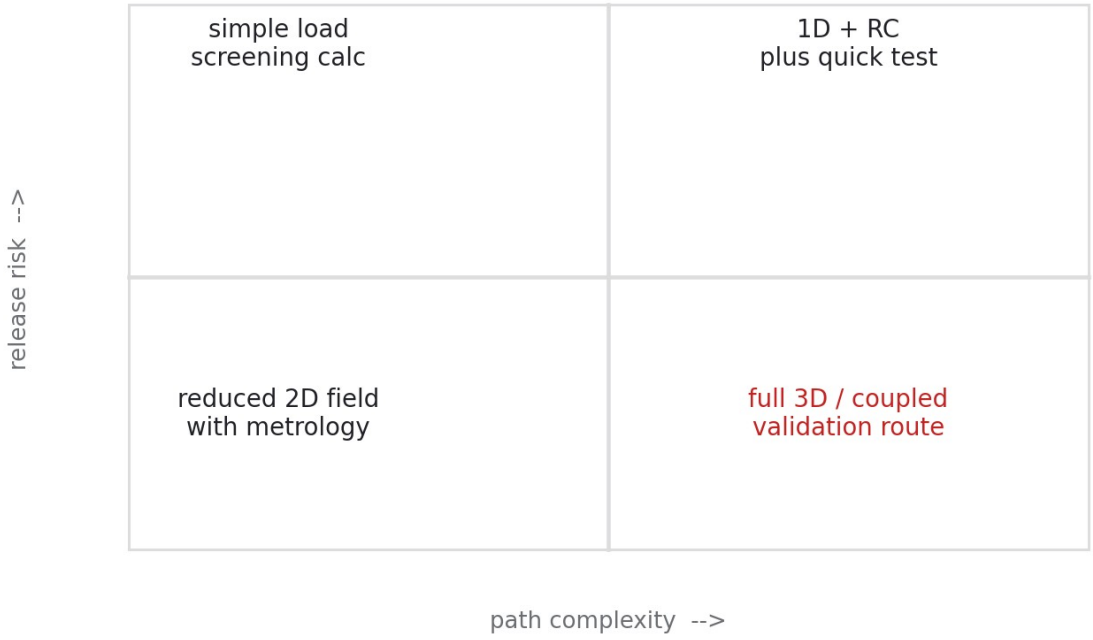


Figure 2: The project-review question is not “which model is nicest” but “which evidence is proportionate to the risk”.

# Contents I

Front matter and primary technical chapters


| | |
|---|---|
| Imprint and document status | 002 |
| Legal and technical notice | 003 |
| Management Summary | 004 |
| Executive Summary - decision system | 005 |
| 1. Problem space, target view, and terminology | 008 |
| 1.1 Reading logic, audiences, and decision questions | 009 |
| 1.2 Multilayer architectures and heat-path types | 010 |
| 1.3 Material windows of representative clad components | 011 |
| 1.4 Reference ranges for materials and boundary data | 012 |
| 2. Steady-state heat distribution in the layer stack | 013 |
| 2.1 Governing equations and energy balance | 014 |
| 2.2 Series resistance and effective properties | 015 |
| 2.3 Layer stack and heat-path rendering | 016 |
| 2.4 Boundary conditions and cooling coupling | 017 |
| 2.5 Temperature profile through thickness | 018 |

# Contents II

Deep-dive chapters, worked examples, and appendices

2.6 Radial conduction in cylindrical components 019
2.7 Lateral spreading and the through-thickness bottleneck 020
2.8 Steady-state design rules 021
3. Transient loads and pulse operation 022
4. Interfaces, bond quality, and anisotropy 031
5. Model hierarchy and regime selection 039
6. Statistical parametrization and robustness 046
7. Technical renderings and application archetypes 057
8. Design rules, governance, and QA 065
9. Worked Example A - steady-state plate 073
10. Worked Example B - pulse loading 076
11. Worked Example C - cylindrical wall 079
Appendix A - formula compendium 082
Appendix B - nomenclature and units 086
Appendix C - representative material windows 087
Appendix D - project checklist and review gates 088
Conclusion and strategic recommendations 089
Contact and document note 090

# 1. Problem space, target view, and terminology

Why heat distribution in clad systems must be read as an architecture problem rather than a single-property problem

The thermal behavior of a clad component emerges from the architecture of the resistance path. Layer order, thickness, directional conductivity, bond state, and coupling to the environment jointly decide whether a system damps, shifts, or amplifies a hot spot. The correct engineering view is therefore architectural: the stack has to be read as an organized path for heat, not as a bag of isolated material values.

This chapter defines the basic terms used throughout the whitepaper. It creates one vocabulary for management, engineering, and quality so that design choices, release arguments, and test plans refer to the same physical quantities and not to parallel narratives.

- Nominal conductivity is an input value, not the system response.
- Cap layer, interlayer, and carrier perform different thermal roles.
- Boundary conditions become part of the material system as soon as the component is operated.
- A useful reading logic does not separate management, engineering, and QA artificially.

# 1.1 Reading logic, audiences, and decision questions

From rapid screening to release-grade sizing

The document is written for three audiences. Management and sales need a robust screening framework, engineering needs a clean translation from load to model choice, and quality needs metrics that can actually be monitored in trials and series production. In multilayer clad materials the temperature field is not set by one conductivity value but by the ordered sequence of layers, interfaces, and release conditions.

A good review asks three questions in sequence: what is the dominant heat path, which variable controls the risk tail, and what level of evidence is proportionate at the current project stage. The practical consequence is that architecture and boundary definition must be reviewed together before a stack is approved or rejected. When those questions are answered in the wrong order, teams tend to over-model low-risk cases and under-test interface-sensitive ones.

- Screening question: is the architecture thermally plausible at all?
- Engineering question: which path and which time scale dominate?
- Release question: which metric must remain stable lot to lot?

# 1.2 Multilayer architectures and heat-path types

Representative stack topologies and dominant path families

A multilayer clad system can be read through several path families: through-thickness conduction, lateral spreading in a conductive cap, radial release in cylindrical walls, and mixed paths that combine local sources with finite sink coupling. In multilayer clad materials the temperature field is not set by one conductivity value but by the ordered sequence of layers, interfaces, and release conditions.

The first design task is not fine optimization but path identification. Once the dominant path is visible, material choice and test strategy become much more efficient because the irrelevant degrees of freedom drop out of the discussion. The practical consequence is that architecture and boundary definition must be reviewed together before a stack is approved or rejected.

- Through-thickness path: typical for walls, liners, and plates under distributed loading.
- Lateral-spreading path: typical for local heat sources and smoothing tasks.
- Radial path: typical for tubes, rolls, and cylindrical process hardware.

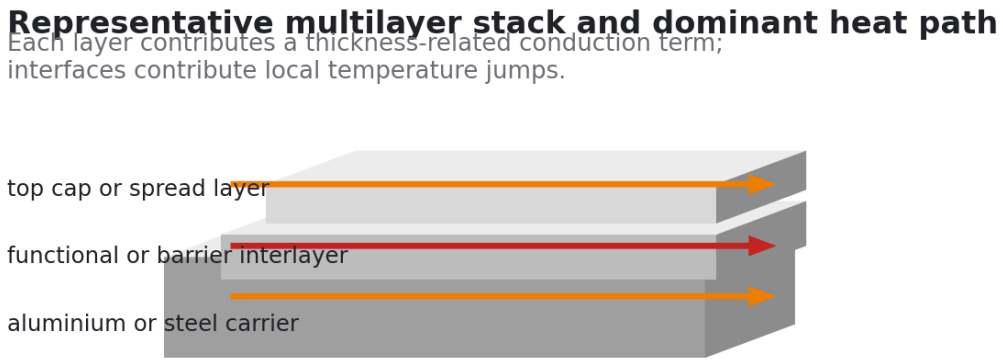


Figure 3: A representative layer stack makes the dominant path legible before numerical work begins.

# 1.3 Material windows of representative clad components

Conductivity classes across metals, interlayers, and barrier systems

Material combinations used in clad products span almost three orders of magnitude in conductivity. Copper and aluminium act as spreaders or low-resistance caps, nickel and stainless systems often act as functional or corrosion-resistant layers, while ceramic or polymer films are thin but thermally expensive. In multilayer clad materials the temperature field is not set by one conductivity value but by the ordered sequence of layers, interfaces, and release conditions.

The relevant question is rarely which material has the highest k in isolation. The useful question is how much thermal duty each layer must carry within the given thickness budget and whether a thin low-k layer is acceptable because it adds function, corrosion resistance, or process stability. The practical consequence is that architecture and boundary definition must be reviewed together before a stack is approved or rejected.

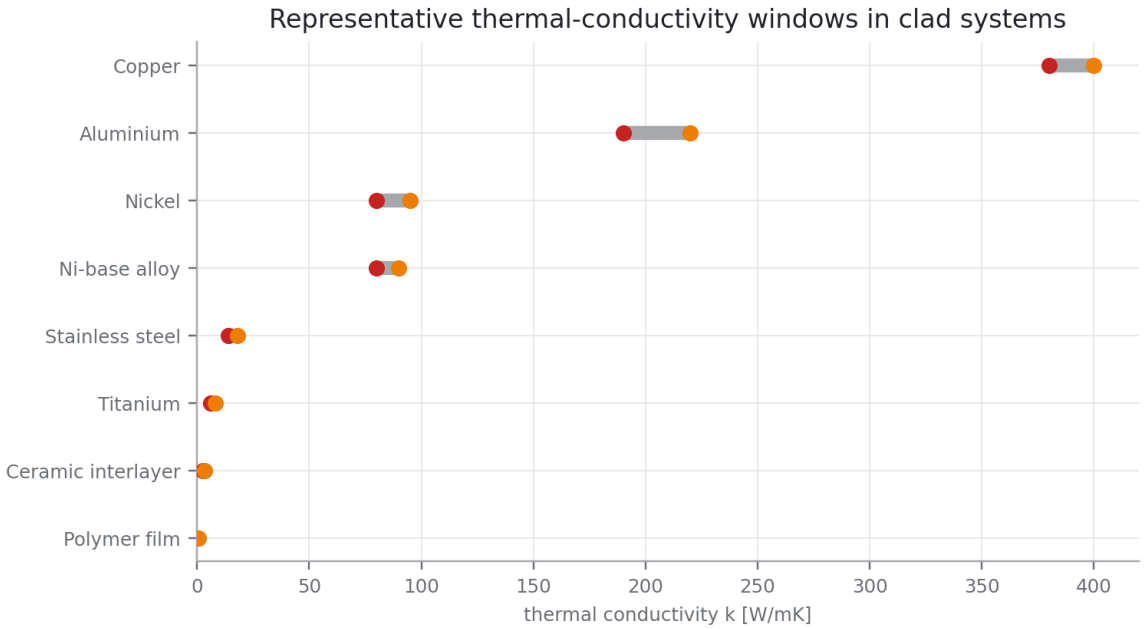


Figure 4: Representative conductivity windows of common metals and functional interlayers in clad architectures.

# 1.4 Reference ranges for materials and boundary data

Typical engineering windows used in this whitepaper

The modeled study behind this whitepaper spans the broad windows typically seen during early design and quotation work. Total wall thickness ranges from approximately 0.26 to 9.65 mm, the cooling coefficient h from about 80 to 3000 W/m²K, total interface resistance from 0.03 to 44 mm²K/W, and pulse duration from roughly 0.15 to 120 s.

These windows are intentionally wider than the final operating point of a validated product. Their purpose is to reveal ranking and sensitivity, not to mimic one exact series design.

| Parameter | Lower reference | Upper reference | Comment |
|---|---|---|---|
| Total thickness L_tot | 0.26 mm | 9.65 mm | Broad industrial screening range |
| Cooling coefficient h | 80 W/m²K | 3000 W/m²K | From weak air-side to strong liquid-side release |
| Total interface resistance ΣR_int | 0.03 mm²K/W | 44 mm²K/W | From dense metallic contact to poor bond state |
| Pulse duration t_p | 0.15 s | 120 s | Short bursts to slow transient duty |
| Representative source flux q'' | 0.02 MW/m² | 0.35 MW/m² | Local process or component loading window |

# 2. Steady-state heat distribution in the layer stack

How the dominant resistance path shapes the stationary field

In steady-state service the task is to understand where the total temperature drop is spent. In a one-dimensional reading every layer and every bond contributes a resistance term, while lateral dimensions and local source size decide whether spreading must be considered in addition.

This chapter translates the stack into a stationary resistance system. The aim is direct: build an engineering argument that identifies the thermal bottleneck before detailed simulation starts.

- Every layer adds a resistance term through thickness.
- Every bond line can inject a local temperature jump.
- Cooling-side mounting and contact quality are part of the same thermal budget.
- Steady-state calculations are the baseline for later hot-spot and robustness review.

# 2.1 Governing equations and energy balance

The stationary field as a balance of conserved heat flow

The stationary starting point is the heat equation without storage term. In the simplest through-thickness case, heat flux is constant and temperature changes only because each layer consumes part of the thermal head. Under steady-state loading the same heat rate crosses every through-thickness layer as long as internal generation is negligible, so the stack can be read as a resistance ledger per unit area.

The engineering value of the balance is that it converts geometry and material data into a transparent resistance statement that can be reviewed by design, sales, and quality alike. This is why early design work should first identify the dominant term in R_th - layer conduction, interface jump, or sink-side release - before adding model complexity.

**Core relations**

```
q" = -k_perp dT/dx
d/dx ( k dT/dx ) = 0
Delta T_total = q" * R_th,total
```

# 2.2 Series resistance and effective properties

Reducing a stack to one equivalent through-thickness property

For fast screening, the multilayer wall can be collapsed into one effective resistance and one effective through-thickness conductivity. This is a powerful reduction because it preserves the quantity that matters most for a first temperature estimate. Under steady-state loading the same heat rate crosses every through-thickness layer as long as internal generation is negligible, so the stack can be read as a resistance ledger per unit area.

The reduction is valid as long as the path really is one-dimensional over the relevant area. Once local loads, edge effects, or severe anisotropy dominate, the equivalent value must be supplemented by field effects. This is why early design work should first identify the dominant term in R_th - layer conduction, interface jump, or sink-side release - before adding model complexity.

**Equivalent-property model**

```
R_th/A = Σ_i ( t_i / k_i ) + Σ_j R_int,j + 1/h
k_eff = L_tot / ( Σ_i ( t_i / k_i ) + Σ_j R_int,j )
```

# 2.3 Layer stack and heat-path rendering

Making the resistance path visible before calculation

A rendering of the stack is often more valuable in early reviews than a dense equation sheet. It forces the team to show which layer spreads heat, which layer provides function, and where the sink is actually connected. Under steady-state loading the same heat rate crosses every through-thickness layer as long as internal generation is negligible, so the stack can be read as a resistance ledger per unit area.

Once the path is drawn, unnecessary arguments disappear quickly. The rendering becomes the bridge between procurement language, design intent, and physical calculation. This is why early design work should first identify the dominant term in R_th - layer conduction, interface jump, or sink-side release - before adding model complexity.

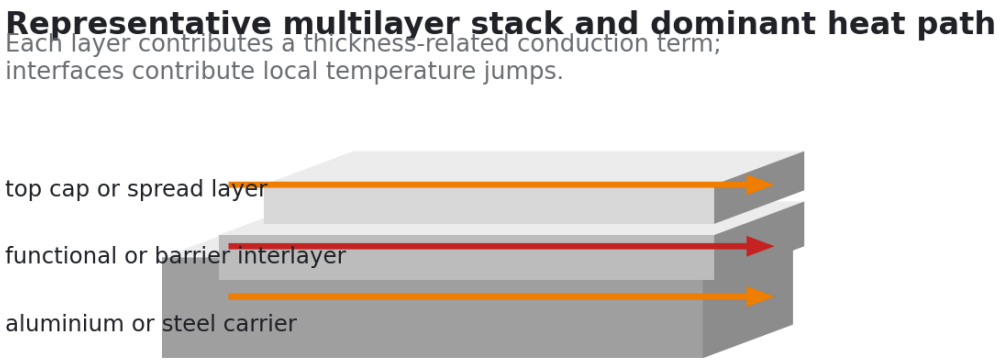


Figure 5: Rendering the stack as a path diagram aligns geometry, function, and thermal responsibility.

# 2.4 Boundary conditions and cooling coupling

Why sink-side definition belongs in the material review

A good clad stack can still fail thermally when the sink-side contact is weak, rough, or variable. The release condition must therefore be treated as part of the same thermal chain as the wall itself. Under steady-state loading the same heat rate crosses every through-thickness layer as long as internal generation is negligible, so the stack can be read as a resistance ledger per unit area.

In practical terms, mounting torque, contact flatness, contact media, oxide layers, and coolant-side evidence can shift the temperature level more than a moderate change in bulk conductivity. This is why early design work should first identify the dominant term in R_th - layer conduction, interface jump, or sink-side release - before adding model complexity.

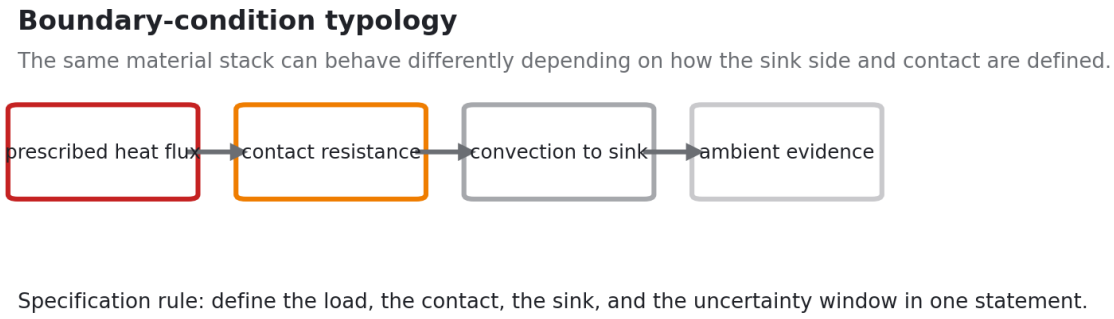


Figure 6: Load, contact, sink, and ambient evidence belong in one boundary-condition statement.

# 2.5 Temperature profile through thickness

Reading the slope and the local jumps

The through-thickness temperature profile is a direct diagnostic tool. A shallow slope indicates low resistance, a steep slope indicates a bottleneck, and a vertical jump indicates interface resistance. Under steady-state loading the same heat rate crosses every through-thickness layer as long as internal generation is negligible, so the stack can be read as a resistance ledger per unit area.

Plotting the profile is not cosmetic. It immediately shows whether the design is dominated by a thin functional layer, by a poorly released carrier, or by the sink-side contact. This is why early design work should first identify the dominant term in R_th - layer conduction, interface jump, or sink-side release - before adding model complexity.

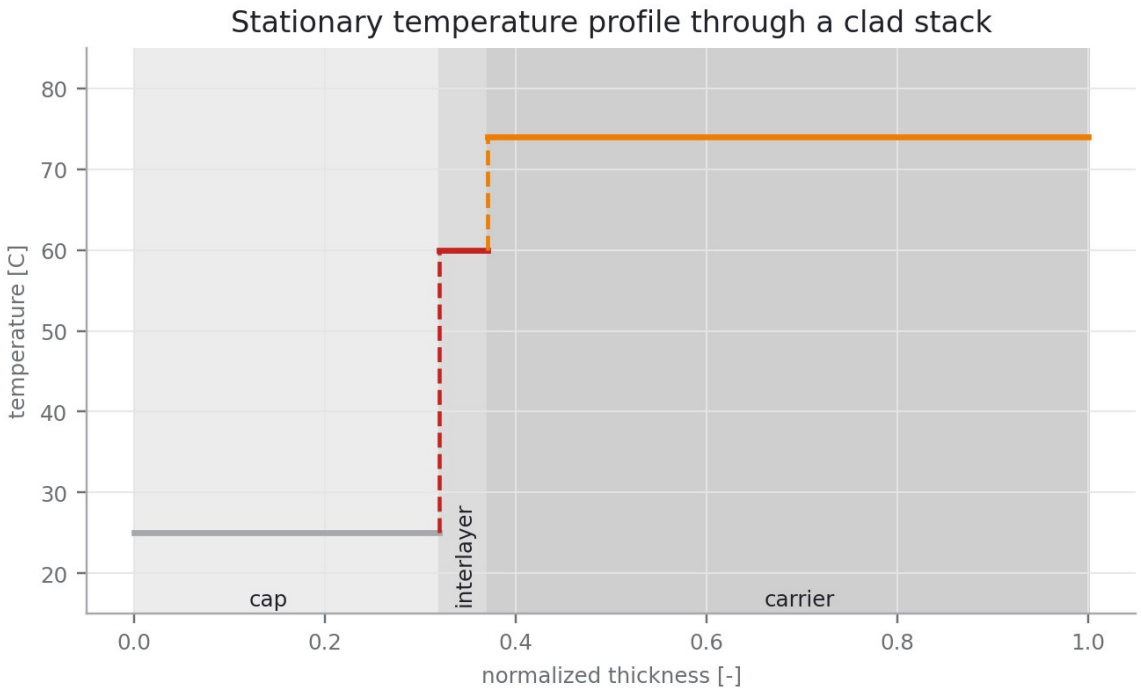


Figure 7: The stationary profile reveals where the thermal head is consumed and where local jumps appear.

# 2.6 Radial conduction in cylindrical components

Why tubes, rolls, and sleeves follow log-shaped profiles

In cylindrical parts the area available for heat flow changes with radius. The result is a logarithmic temperature distribution rather than the linear profile known from planar walls. Under steady-state loading the same heat rate crosses every through-thickness layer as long as internal generation is negligible, so the stack can be read as a resistance ledger per unit area.

This matters for clad tubes and rolls because a thin high-resistance layer near the inner or outer surface can influence the temperature field more strongly than its thickness alone would suggest in a flat-wall approximation. This is why early design work should first identify the dominant term in R_th - layer conduction, interface jump, or sink-side release - before adding model complexity.

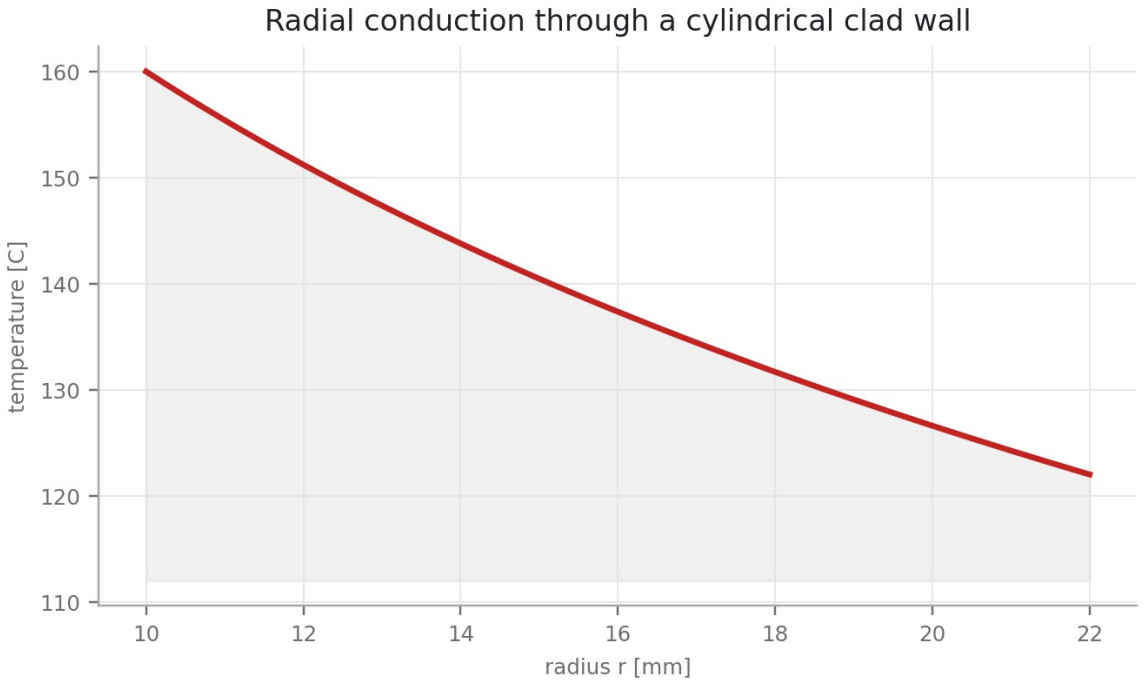


Figure 8: Radial conduction must respect changing area and therefore uses a logarithmic profile.

# 2.7 Lateral spreading and the through-thickness bottleneck

Two resistance families acting together

Many real components combine a local source with finite sink coupling. In those cases the top layer first has to spread the load laterally before the wall can release it through thickness. Under steady-state loading the same heat rate crosses every through-thickness layer as long as internal generation is negligible, so the stack can be read as a resistance ledger per unit area.

A highly conductive cap reduces local gradients, but the field will still collapse into a hot spot when the through-thickness leg beneath it is too resistive. Design should therefore separate spreading benefit from release capacity. This is why early design work should first identify the dominant term in R_th - layer conduction, interface jump, or sink-side release - before adding model complexity.

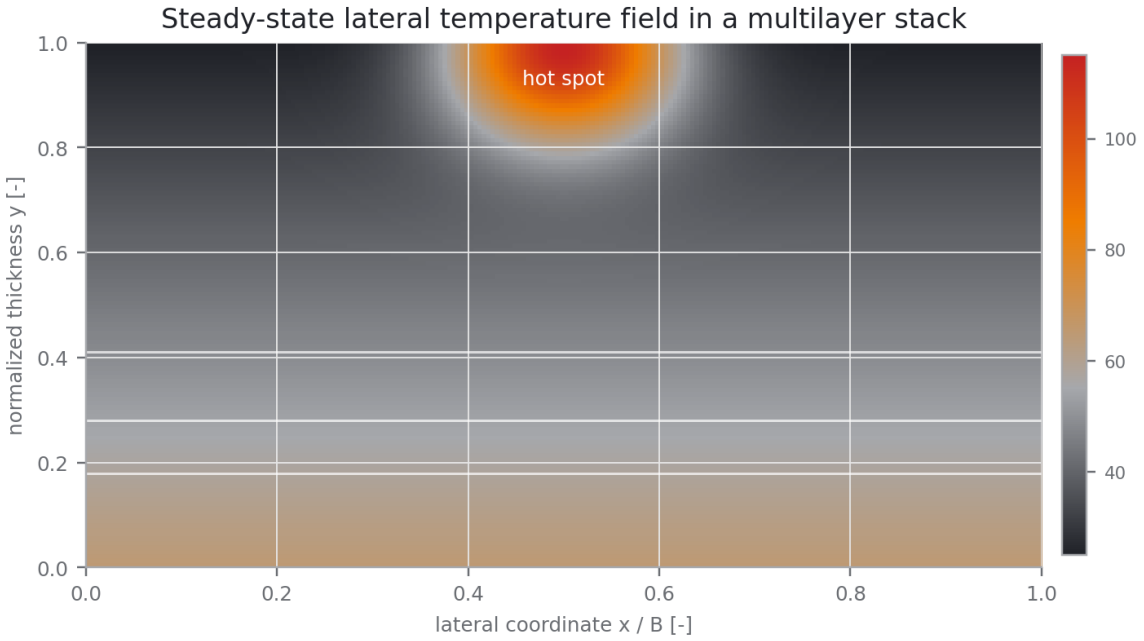


Figure 9: A local hot spot can remain despite a conductive cap when the through-thickness path below it is bottlenecked.

# 2.8 Steady-state design rules

Direct heuristics for early sizing

Steady-state design can be made surprisingly robust with a few disciplined heuristics. Put conductivity where spreading is needed, keep thermally expensive interlayers thin, and define the sink-side evidence at the same time as the material stack. Under steady-state loading the same heat rate crosses every through-thickness layer as long as internal generation is negligible, so the stack can be read as a resistance ledger per unit area.

Do not optimize the nominal point alone. The better design is the one whose hot-spot temperature remains acceptable across manufacturing scatter, mounting variation, and realistic release conditions. This is why early design work should first identify the dominant term in R_th - layer conduction, interface jump, or sink-side release - before adding model complexity.

- Prefer thin functional layers with explicit justification for every low-k addition.
- Treat ΣR_int and 1/h as first-class design terms, not as afterthoughts.
- Use field models only after the dominant steady-state path has been made visible.

# 3. Transient loads and pulse operation

Time scale, penetration depth, and stored thermal mass

A stack that looks harmless in steady state can still fail under pulsed service. Short loads interact with thermal penetration depth and nodal storage, not just with total resistance. The correct question becomes: how deep does the pulse reach within the available time, and which nodes can store or release energy quickly enough?

This chapter introduces the transient logic used throughout the whitepaper. It links unsteady conduction, reduced RC models, and measurement windows so that pulse-dominated applications can be classified without jumping immediately to a full multiphysics model.

- The decisive quantity is the ratio between pulse duration and thermal response time.
- Short pulses probe only part of the stack; long pulses activate the full resistance ledger.
- Transient design requires both resistance and thermal mass.

# 3.1 Unsteady equation and initial conditions

Storage enters the balance as soon as temperature changes with time

The unsteady heat equation adds a storage term to the stationary balance. This term is what turns a wall into a dynamic object whose response depends on both conductivity and volumetric heat capacity. In transient service the relevant question is not only how large the thermal resistance is, but how fast the layer stack stores and redistributes energy.

For project work, initial condition discipline matters. The same pulse has a different consequence on a cold stack, on a preheated stack, and on a stack that has not fully recovered from the previous cycle. Engineering decisions should therefore be tied to the characteristic time scale of the load, the penetration depth, and the homogenization time t95 rather than to k alone.

**Unsteady relations**

```
∂T/∂t = α ∂²T/∂x²
α = k / (ρ c_p)
Initial state: T(x,0) = T0(x)
```

# 3.2 Thermal penetration depth

How far a transient event reaches into the wall

Thermal penetration depth is a compact way to judge whether a pulse acts only in the cap, reaches the carrier, or engages the full wall. It is one of the most useful screening quantities for transiently loaded clad components. In transient service the relevant question is not only how large the thermal resistance is, but how fast the layer stack stores and redistributes energy.

When the penetration depth remains below the total wall thickness, local layers dominate the response. When it exceeds the wall thickness, the full stack and the sink-side coupling come into play. Engineering decisions should therefore be tied to the characteristic time scale of the load, the penetration depth, and the homogenization time t95 rather than to k alone.

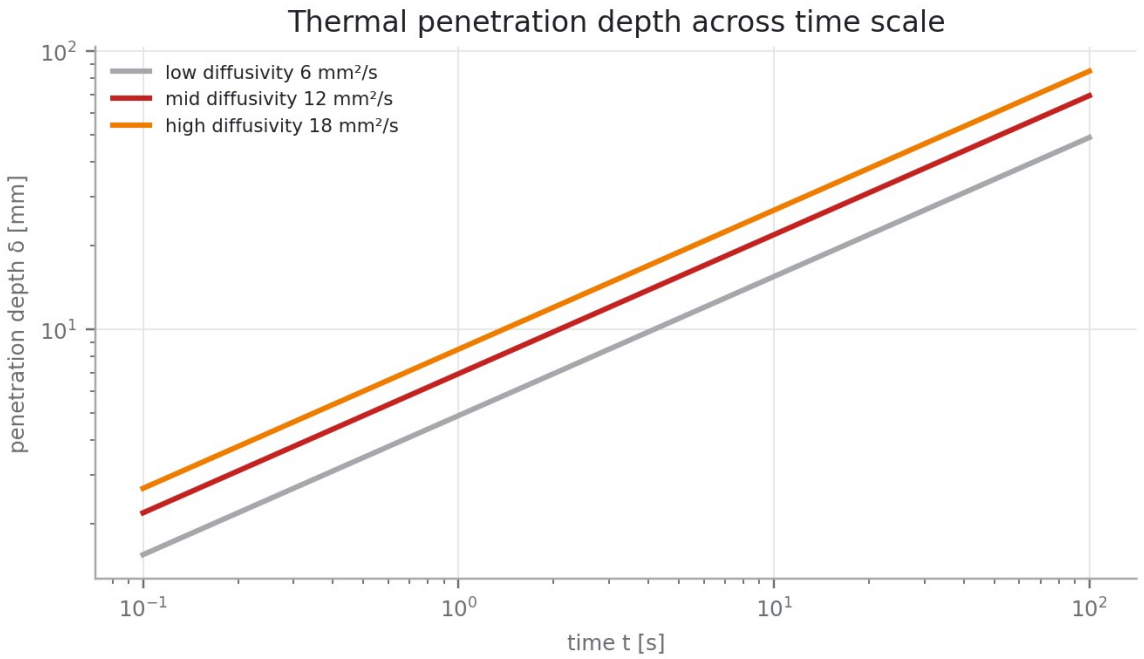


Figure 10: Penetration depth grows with the square root of time and with thermal diffusivity.

# 3.3 Time constants, eigenmodes, and t95

Reading how quickly a stack approaches thermal equilibrium

Every multilayer system has characteristic response times. In reduced models they appear as time constants; in field models they appear as modal decay behavior. Both describe how fast the temperature field forgets an initial disturbance. In transient service the relevant question is not only how large the thermal resistance is, but how fast the layer stack stores and redistributes energy.

The practical metric used in this document is t95, the time required for 95 percent of the final temperature response or homogenization level. It is easy to communicate and maps well to cyclic duty decisions. Engineering decisions should therefore be tied to the characteristic time scale of the load, the penetration depth, and the homogenization time t95 rather than to k alone.

**Characteristic-time view**

```
τ ≈ R_th * C_th
t95 ≈ 3 τ
Fo = α t / L²
```

# 3.4 Reduced discretization and RC/FVM approach

Fast dynamic models for design iteration

A reduced RC or finite-volume model is often the best compromise between speed and physical fidelity. It resolves storage per layer, preserves interface terms, and is still cheap enough for parameter sweeps and robust design studies. In transient service the relevant question is not only how large the thermal resistance is, but how fast the layer stack stores and redistributes energy.

The value of this model class is practical: it allows engineering teams to run many cases, compare route options, and calibrate against measurement without the cost of a full three-dimensional transient simulation. Engineering decisions should therefore be tied to the characteristic time scale of the load, the penetration depth, and the homogenization time t95 rather than to k alone.

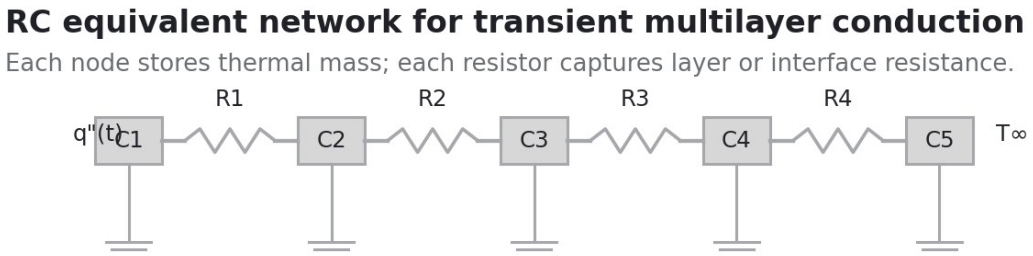


Figure 11: RC nodes store thermal mass, while resistors represent layer and interface impedances.

# 3.5 Transient temperature profiles in the layer stack

Why the profile changes shape with time

During a pulse the temperature profile is not fixed. Early in time only the region near the source responds, while deeper layers remain almost unchanged. As time grows, the field penetrates further and the profile gradually approaches its steady-state shape. In transient service the relevant question is not only how large the thermal resistance is, but how fast the layer stack stores and redistributes energy.

This evolution is essential for interpreting sensor positions. A measurement close to the source may show a sharp peak, while a sensor near the sink may barely react during the same event. Engineering decisions should therefore be tied to the characteristic time scale of the load, the penetration depth, and the homogenization time t95 rather than to k alone.

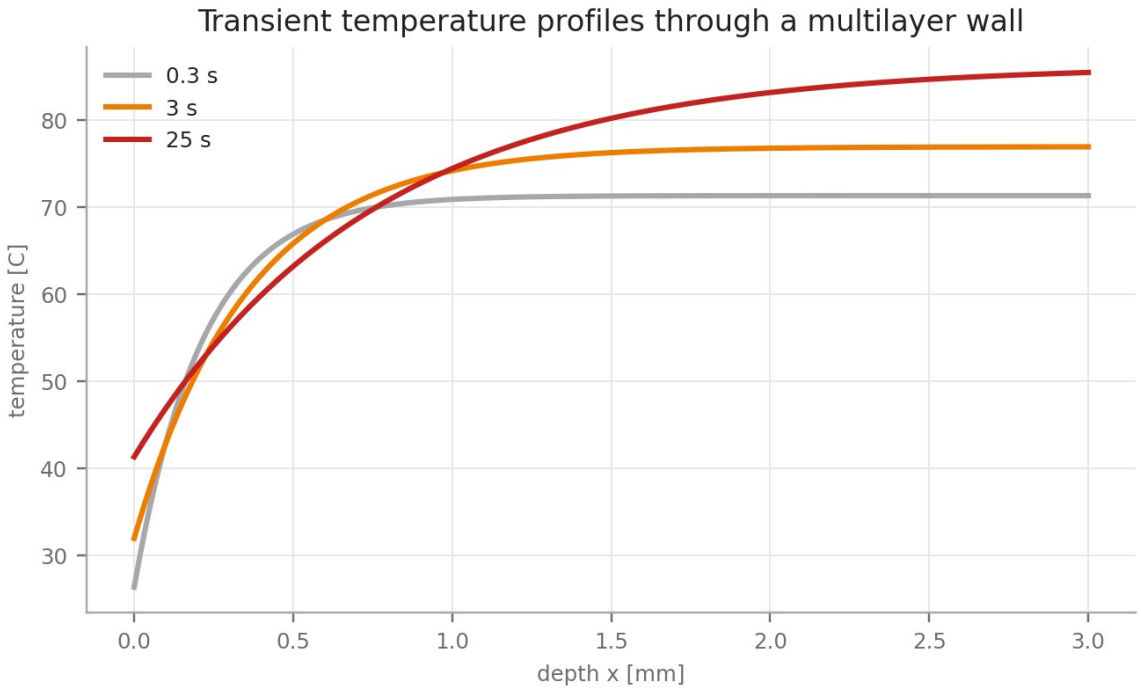


Figure 12: Early-time profiles are shallow in depth; later profiles engage more of the wall.

# 3.6 Load and time parameters for pulse assessments

A compact set of transient descriptors

Pulse analysis becomes manageable when a small set of descriptors is tracked consistently: pulse height, pulse duration, repetition rate, duty factor, initial temperature, and sink-side response. These variables capture most of the relevant decision logic for early transient reviews. In transient service the relevant question is not only how large the thermal resistance is, but how fast the layer stack stores and redistributes energy.

The objective is not to create a large parameter inventory. The objective is to isolate the variables that move the thermal peak, the recovery time, and the lot-to-lot sensitivity. Engineering decisions should therefore be tied to the characteristic time scale of the load, the penetration depth, and the homogenization time t95 rather than to k alone.

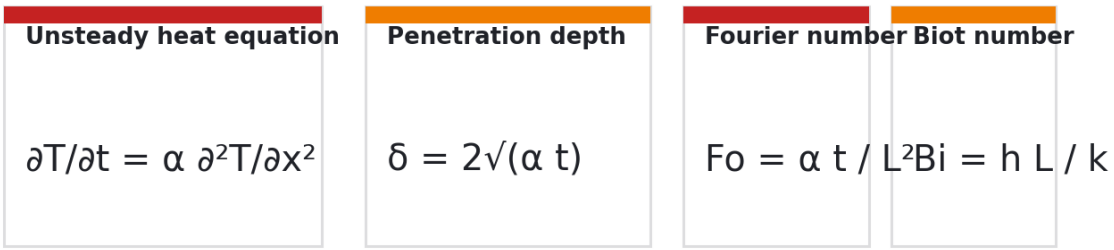


Figure 13: A small family of dimensionless and physical descriptors is enough for most pulse-screening tasks.

# 3.7 Explicit stability, implicit robustness, and computational economy

Choosing a time-integration scheme that fits the task

Explicit transient schemes are simple and fast for very small time steps, but they are limited by stability criteria. Implicit schemes are more robust and usually preferable when long transients, stiff interfaces, or production-grade sweeps are part of the workflow. In transient service the relevant question is not only how large the thermal resistance is, but how fast the layer stack stores and redistributes energy.

The rule is straightforward: choose the method that gives the necessary evidence at the lowest computational cost. In many industrial RC and 1D FVM studies, implicit integration is the better operational default. Engineering decisions should therefore be tied to the characteristic time scale of the load, the penetration depth, and the homogenization time t95 rather than to k alone.

- Explicit: good for simple short transients and educational models.
- Implicit: robust for stiff stacks, long pulses, and parameter sweeps.
- Model economy matters because robust design requires many evaluated cases.

# 3.8 Measurement windows and validation strategy

Sensors only help when they look at the right time scale

Transient validation is as much about timing as about absolute temperature. A slow sensor may miss the local peak entirely, while an infrared snapshot without emissivity control may misread the field shape. In transient service the relevant question is not only how large the thermal resistance is, but how fast the layer stack stores and redistributes energy.

The validation strategy should therefore define where the thermal signal is expected, how fast it develops, and which sensor bandwidth is required. This is especially important when the design argument rests on t95 or on short-lived interface hot spots. Engineering decisions should therefore be tied to the characteristic time scale of the load, the penetration depth, and the homogenization time t95 rather than to k alone.

- Place sensors where the dominant path predicts the strongest signal.
- Define sensor bandwidth against the expected transient rise time.
- Use one measurement route for absolute level and another for field shape when necessary.

# 4. Interfaces, bond quality, and anisotropy

Why bonds, oxide films, and directional microstructure govern local gradients

Interfaces are where multilayer clad systems stop behaving like textbook bulk materials. Contact quality, bonding route, roughness, oxide films, and diffusion zones determine whether heat crosses smoothly or through a locally resistive bridge.

This chapter treats the interface as a first-class thermal feature. It also extends the discussion to anisotropy, because process history and directional microstructure often create different conductivities parallel and perpendicular to the layer plane.

- Interface resistance appears as a local temperature jump.
- Micro-contact statistics matter more than nominal thickness alone.
- Directional microstructure can improve spreading while penalizing through-thickness equalization.

# 4.1 Interface resistance as temperature jump

The shortest local equation in the whitepaper is often the most important one

Interface resistance acts as a localized thermal discontinuity. The temperature jump across the bond is directly proportional to heat flux and to the effective resistance of the interface zone. In practice the bond line often controls local temperature gradients more strongly than the nominal conductivity of the adjacent bulk materials.

Because the jump is local, it is easy to underestimate in global averages. Yet exactly this local jump often determines whether a coating, liner, or cap sees an unacceptable peak temperature. Interface effects should be treated as release-critical parameters whenever hot spots, cyclic loads, or uncertain contact quality are part of the duty profile.

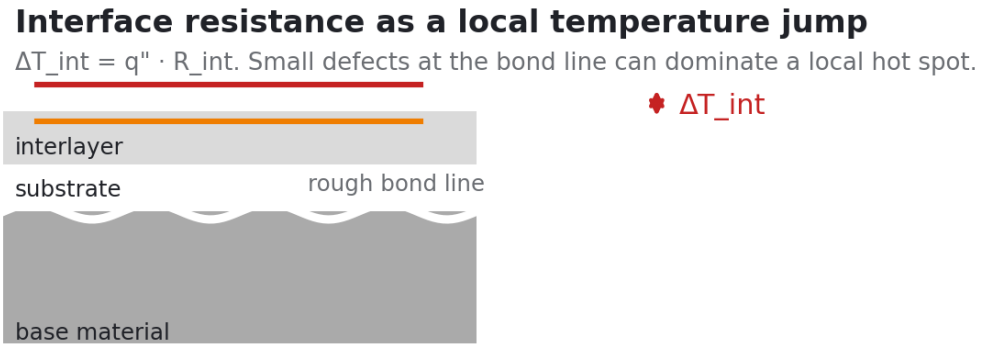


Figure 14: A bond-line temperature jump may dominate the local hot-spot signature even when bulk k looks favorable.

# 4.2 Microstructure, oxide films, and micro-contacts

What sits inside a “bond” from a thermal point of view

From a thermal perspective a bond is a statistical ensemble of contact bridges, diffusion regions, oxides, pores, and roughness-driven gaps. The local effective resistance depends on how these features connect into a real heat path. In practice the bond line often controls local temperature gradients more strongly than the nominal conductivity of the adjacent bulk materials.

This is why nominal process labels alone are insufficient. Two stacks built from the same material pair can behave differently when oxide removal, pressure history, or local micro-contact density differ. Interface effects should be treated as release-critical parameters whenever hot spots, cyclic loads, or uncertain contact quality are part of the duty profile.

- Oxide films add resistance despite very small thickness.
- True metallic bridges carry most of the local heat transfer.
- Porosity and lack of continuity broaden the risk tail even when the mean looks acceptable.

# 4.3 Statistical scatter of bond quality

Why production capability must be read as a thermal variable

Bond quality never exists as one deterministic value in production. Real series behavior shows a distribution that can be narrow and dense, moderately dispersed, or strongly tailed when porosity and local separation occur. In practice the bond line often controls local temperature gradients more strongly than the nominal conductivity of the adjacent bulk materials.

For release this means that thermal approval should not rely on one representative sample only. The relevant question is how much the bond scatter broadens the peak-temperature distribution of the finished part. Interface effects should be treated as release-critical parameters whenever hot spots, cyclic loads, or uncertain contact quality are part of the duty profile.

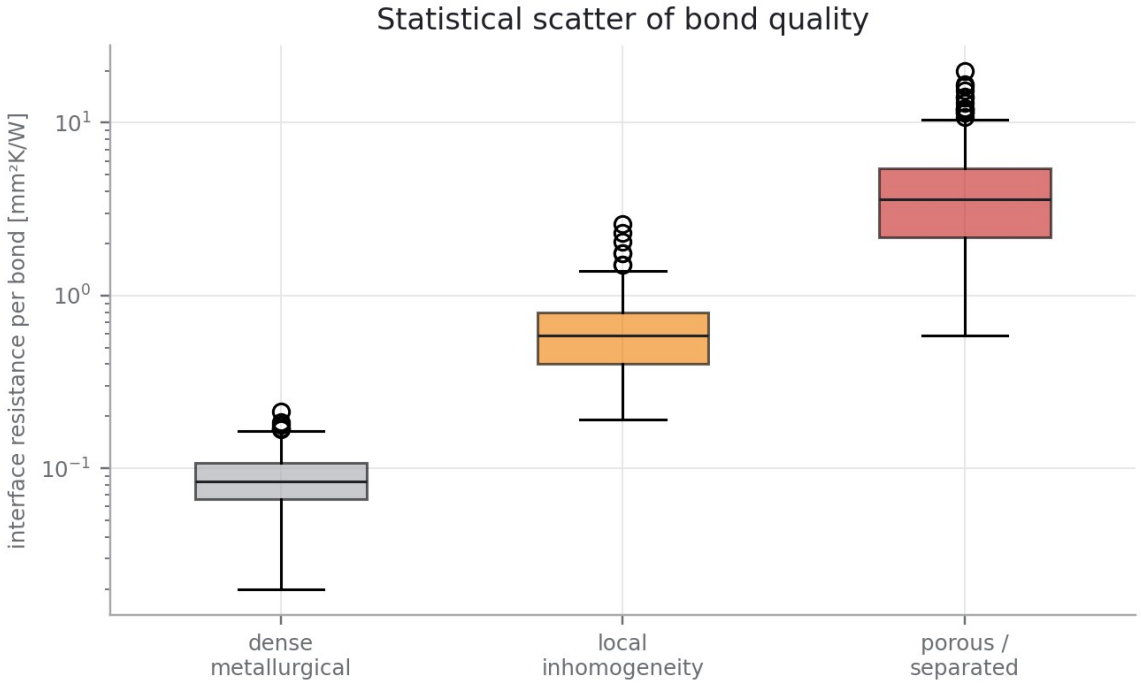


Figure 15: Representative bond classes show markedly different distributions of interface resistance.

# 4.4 Sensitivity of the field to ΣR_int

How total interface resistance shifts the hot-spot level

The sum of interface resistances is a compact variable for quick sensitivity review. Increasing ΣR_int shifts the local peak upward, and the slope of that shift increases with heat flux. In practice the bond line often controls local temperature gradients more strongly than the nominal conductivity of the adjacent bulk materials.

The implication is direct: high-flux service magnifies otherwise moderate bond imperfections. Interface quality therefore becomes progressively more important as duty becomes sharper or more concentrated. Interface effects should be treated as release-critical parameters whenever hot spots, cyclic loads, or uncertain contact quality are part of the duty profile.

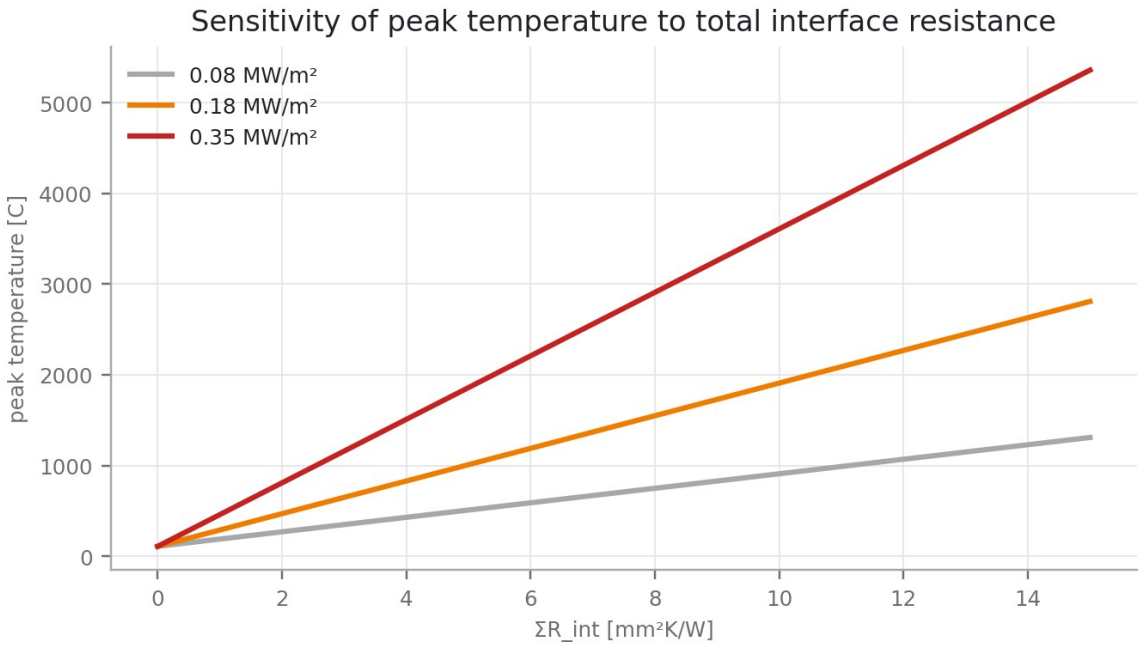


Figure 16: Higher heat flux makes the temperature penalty of ΣR_int progressively steeper.

# 4.5 Roughness, pressure, and chemical boundary conditions

The interface is influenced by assembly and service chemistry

Surface roughness and assembly pressure decide how much true contact area is created at the bond or at a mounted sink interface. Chemical service can then alter the picture over time by oxidation, contamination, or corrosion-assisted degradation. In practice the bond line often controls local temperature gradients more strongly than the nominal conductivity of the adjacent bulk materials.

Good design therefore separates three states of the same interface: as-manufactured, as-mounted, and as-operated. The release logic should state which of these conditions is represented by the supporting evidence. Interface effects should be treated as release-critical parameters whenever hot spots, cyclic loads, or uncertain contact quality are part of the duty profile.

- Roughness changes real contact area and therefore local conductance.
- Pressure can improve contact, but only within the mechanically admissible window.
- Chemical exposure may shift the interface state over life and must be part of durability review.

# 4.6 Anisotropy and directed microstructure

Different conductivities in different directions

Rolled, forged, or directionally solidified structures often conduct heat better along the layer plane than across it. In clad systems this can be beneficial for smoothing but detrimental for rapid through-thickness release. In practice the bond line often controls local temperature gradients more strongly than the nominal conductivity of the adjacent bulk materials.

The correct engineering move is not to reject anisotropy by default. It is to decide whether the application primarily needs spreading, penetration, or a controlled compromise between both. Interface effects should be treated as release-critical parameters whenever hot spots, cyclic loads, or uncertain contact quality are part of the duty profile.

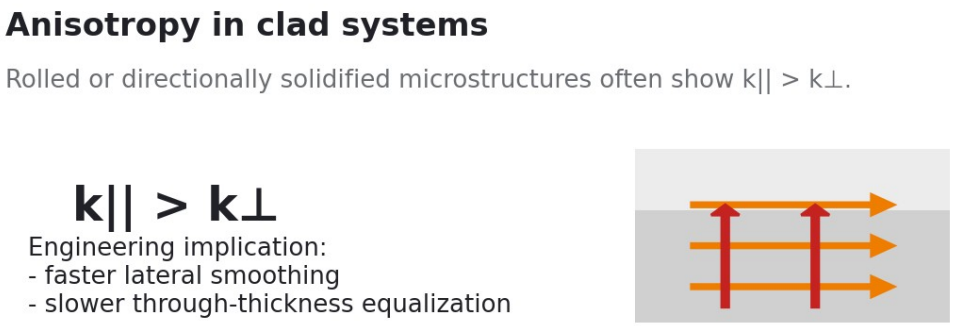


Figure 17: Directional microstructure often increases lateral smoothing while slowing through-thickness equalization.

# 4.7 Process routes and thermal signatures

Manufacturing route as part of the thermal fingerprint

Roll bonding, diffusion bonding, laser cladding, overlay processes, and explosive bonding leave different microstructural and interface signatures behind. These signatures influence field uniformity, directional behavior, and the scatter band of the final thermal response. In practice the bond line often controls local temperature gradients more strongly than the nominal conductivity of the adjacent bulk materials.

For commercial projects this is valuable because process choice can be argued not only on cost or corrosion performance but also on the expected thermal signature and the release effort attached to it. Interface effects should be treated as release-critical parameters whenever hot spots, cyclic loads, or uncertain contact quality are part of the duty profile.

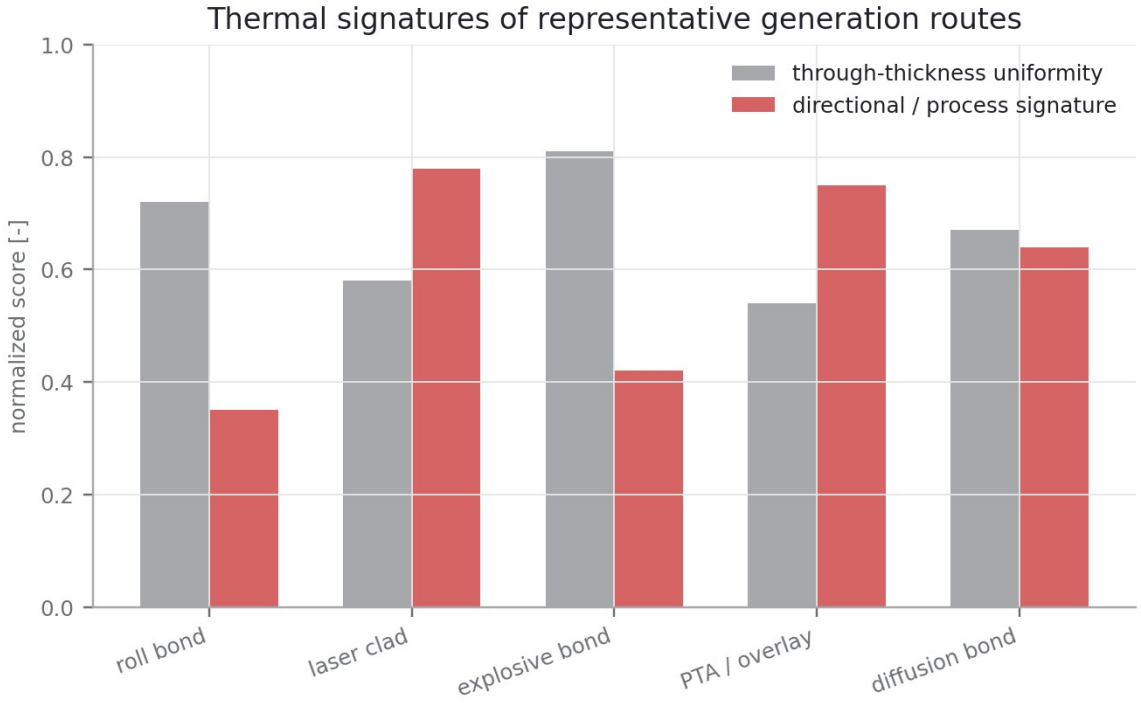


Figure 18: Thermal uniformity and directional process signature vary across representative generation routes.

# 5. Model hierarchy and regime selection

How to choose the right model for the right decision

Thermal modeling is part of engineering governance. The wrong model wastes time or creates false confidence; the right model answers the required question with proportionate effort. The key is not software sophistication but regime fit.

This chapter provides the selection logic used in the whitepaper. It explains when 1D is sufficient, when 2D spreading must be resolved, when 3D geometry becomes unavoidable, and how analytical, reduced, and high-fidelity models can be combined in a staged review route.

- Start with the question, not with the software.
- Escalate model depth when the physics or the release risk demands it.
- Keep reduced models alive even when a field model exists; they are essential for sensitivity work.

# 5.1 Analytical models as governance tools

Why simple equations still matter in complex projects

Analytical models are fast, transparent, and audit-friendly. They expose which term dominates and give the review board a physical explanation that remains visible even after a larger simulation is introduced. The model is not an academic choice; it is a governance tool that determines how much evidence is needed and how quickly decisions can be made.

Used correctly, they are not a primitive substitute for simulation. They are the front-end governance layer that decides whether more modeling is needed at all and where the high-fidelity effort should focus. Good engineering uses the simplest model that still resolves the controlling physics, then escalates only when geometry, interfaces, or uncertainty justify it.

**Governance equations**

```
R_th/A = Σ(t_i/k_i) + ΣR_int + 1/h
τ ≈ R_th * C_th
Screening decisions use these first-order quantities as a risk filter.
```

# 5.2 Criteria for 1D, 2D, and 3D

A practical regime-selection checklist

One-dimensional models are appropriate when the dominant path is through thickness over an area that is large relative to wall thickness. Two-dimensional models become necessary when local sources, edge effects, or finite spreading layers matter. Three-dimensional models are reserved for strongly non-separable geometry, contact patch complexity, or coupled multiphysics. The model is not an academic choice; it is a governance tool that determines how much evidence is needed and how quickly decisions can be made.

In review practice the regime can often be identified from three questions: is the source local, is the geometry laterally bounded, and do multiple release paths compete? Good engineering uses the simplest model that still resolves the controlling physics, then escalates only when geometry, interfaces, or uncertainty justify it.

- 1D: distributed loading and simple release geometry.
- 2D: local source plus finite spreading distance or edge influence.
- 3D: complex contact patches, local features, or non-uniform release on multiple sides.

# 5.3 Biot-Fourier regime map

Using dimensionless groups to place the task

Biot and Fourier numbers provide a compact way to place an application in regime space. Biot compares internal conduction with boundary-side release; Fourier compares elapsed time with the characteristic diffusion time. The model is not an academic choice; it is a governance tool that determines how much evidence is needed and how quickly decisions can be made.

Plotting the operating point does not solve the problem by itself, but it quickly shows whether lumped, thickness-dominated, or full field behavior is likely to control the answer. Good engineering uses the simplest model that still resolves the controlling physics, then escalates only when geometry, interfaces, or uncertainty justify it.

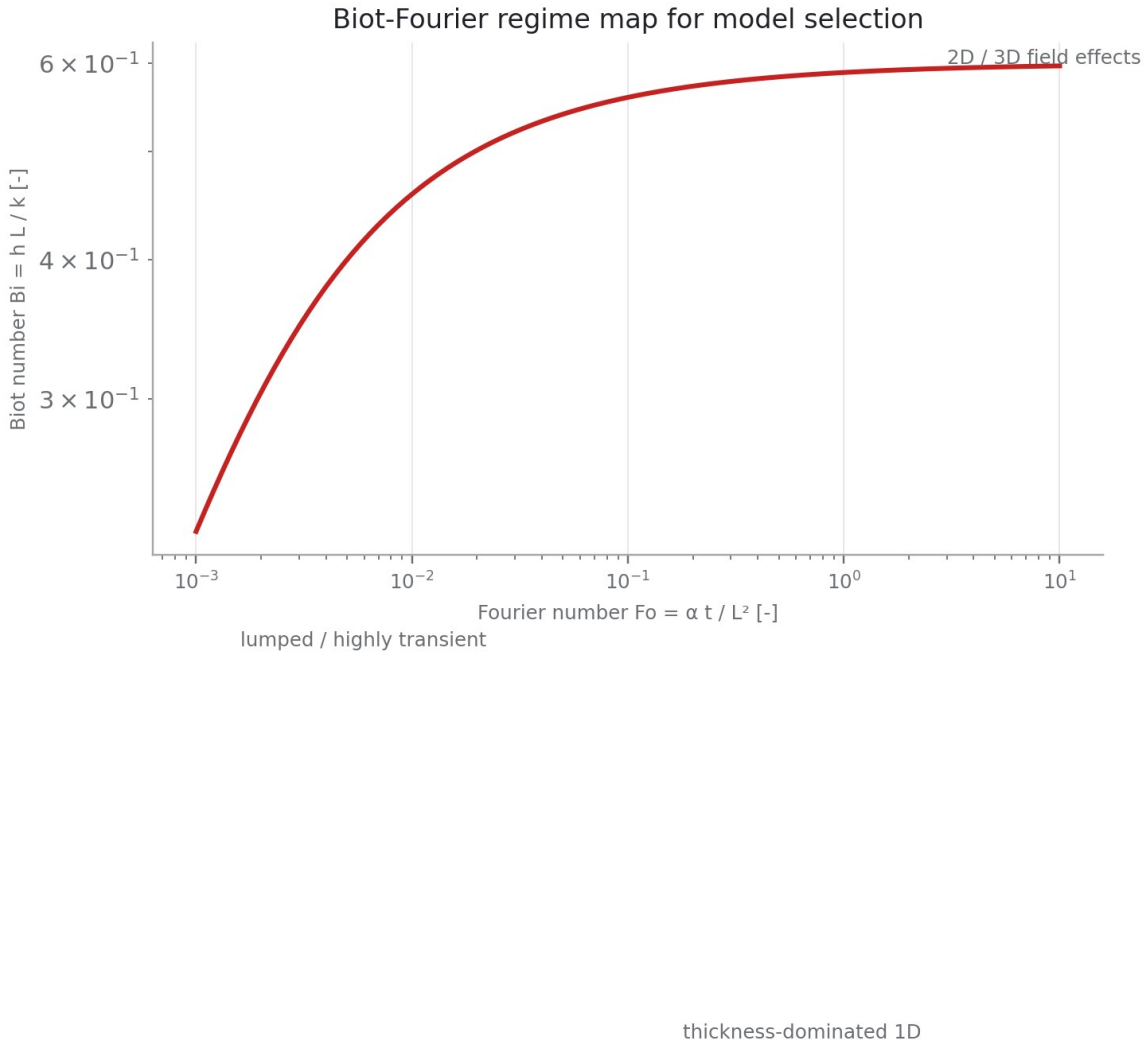


Figure 19: The Biot-Fourier map guides model selection by time scale and release dominance.

# 5.4 Response quantities and homogenization metrics

Choosing the metric that matches the business question

Different projects require different thermal outputs. Some need the absolute peak temperature, some need the temperature spread across a plate, some need t95, and some need a case-to-case transfer metric for production QA. The model is not an academic choice; it is a governance tool that determines how much evidence is needed and how quickly decisions can be made.

The critical governance step is to make one response quantity primary. Once the primary metric is fixed, model calibration, test design, and specification language become much more coherent. Good engineering uses the simplest model that still resolves the controlling physics, then escalates only when geometry, interfaces, or uncertainty justify it.

- Peak temperature for damage or safety review.
- Temperature spread for product quality and homogenization tasks.
- t95 for cyclic duty and start-up behavior.
- Transfer or drift metric for production monitoring.

# 5.5 Model hierarchy in the engineering process

From hand calculation to validated digital twin

Model hierarchy is not a ladder that discards earlier steps. It is a layered toolkit. Hand calculation supports plausibility, RC or 1D FVM supports parametric work, 2D/3D models support local geometry, and the validated digital twin integrates all evidence for late-stage release and optimization. The model is not an academic choice; it is a governance tool that determines how much evidence is needed and how quickly decisions can be made.

Keeping these model layers connected is economically important because each layer answers a different question at a different speed. Good engineering uses the simplest model that still resolves the controlling physics, then escalates only when geometry, interfaces, or uncertainty justify it.

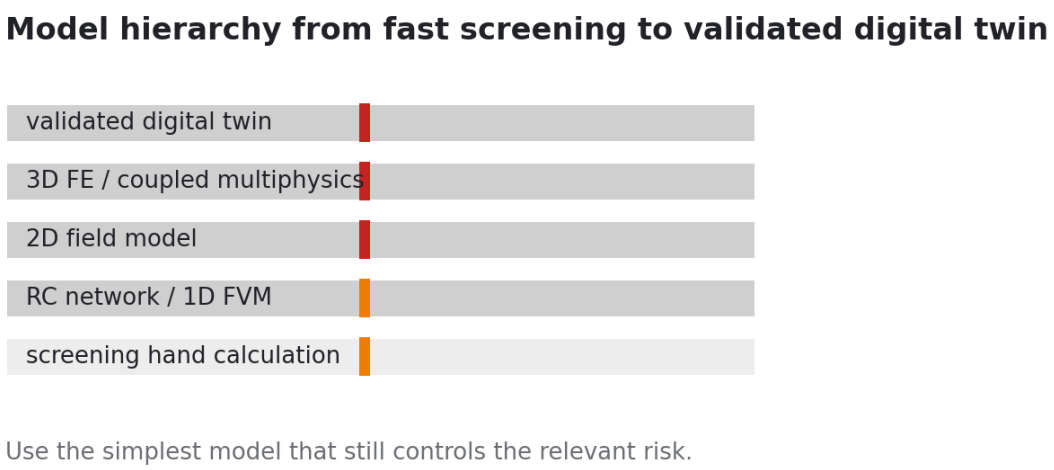


Figure 20: A robust workflow preserves simple models even when more detailed ones are added later.

# 5.6 Calibration, inverse identification, and uncertainty

Turning model output into evidence

Thermal models become operational when they are calibrated against measurement. Inverse identification is often required because interface resistance, contact quality, and release coefficients are not known precisely a priori. The model is not an academic choice; it is a governance tool that determines how much evidence is needed and how quickly decisions can be made.

Calibration should not hide uncertainty. The correct output is a calibrated parameter window together with the residual uncertainty that remains relevant for release. Good engineering uses the simplest model that still resolves the controlling physics, then escalates only when geometry, interfaces, or uncertainty justify it.

- Use measurement to identify uncertain interfaces and release coefficients.
- Keep the calibrated range visible; do not compress it prematurely into one nominal value.
- Translate residual uncertainty into peak-temperature and t95 uncertainty bands.

# 6. Statistical parametrization and robustness

A modeled study of virtual clad stacks

The statistical chapter summarizes 12,000 virtual stacks generated within the reference windows of this whitepaper. Each virtual stack combines thickness, material window, interface condition, release coefficient, and pulse duration in order to show how the distribution of thermal response behaves beyond one nominal point.

The aim is robustness thinking. A design that looks strong in one parameter set can still be weak in the tail. The following pages therefore focus on distributions, sensitivity ranking, and operating windows rather than on one deterministic example alone.

- Modeled samples: 12,000
- Layer count: 2 to 5
- Thickness window: 0.26 to 9.65 mm
- Cooling window h: 80 to 3000 W/m²K
- Pulse window: 0.15 to 120 s

# 6.1 Study design and parameter space

How the virtual stacks were generated

The virtual population combines representative conductive caps, structural carriers, and optional functional interlayers. Thickness is drawn over a wide industrial window, interface resistance is allowed to vary from dense metallic to poor-contact behavior, and boundary-side release spans weak to strong cooling. The statistical study in this whitepaper is based on 12,000 virtual clad stacks spanning thickness, conductivity, interface, cooling, and pulse windows typical of industrial screening work.

This setup intentionally exaggerates the design space relative to one final product. That is useful because a broad study reveals variable ranking and non-linear interactions more clearly than a narrowly tuned product model. The purpose is not to mimic one exact product but to show which variables move the risk distribution and which ones mainly influence the nominal point.

| Variable | Distribution logic | Engineering window |
|---|---|---|
| Layer count | Discrete uniform | 2 to 5 |
| Total thickness | Log-uniform | 0.26 to 9.65 mm |
| Cooling coefficient h | Log-uniform | 80 to 3000 W/m²K |
| Total ΣR_int | Long-tailed log window | 0.03 to 44 mm²K/W |
| Pulse duration | Log-uniform | 0.15 to 120 s |
| Representative q" | Triangular screening load | 0.02 to 0.35 MW/m² |

# 6.2 Distribution of k_eff

Effective through-thickness conductivity across the study

The modeled distribution of effective through-thickness conductivity is broad. The median of the study is approximately 19.5 W/mK, while the 90th percentile reaches about 87.6 W/mK. The statistical study in this whitepaper is based on 12,000 virtual clad stacks spanning thickness, conductivity, interface, cooling, and pulse windows typical of industrial screening work.

This spread confirms a central message of the whitepaper: geometry, thin functional layers, and interfaces can collapse the effective property even when a highly conductive cap is present. The purpose is not to mimic one exact product but to show which variables move the risk distribution and which ones mainly influence the nominal point.

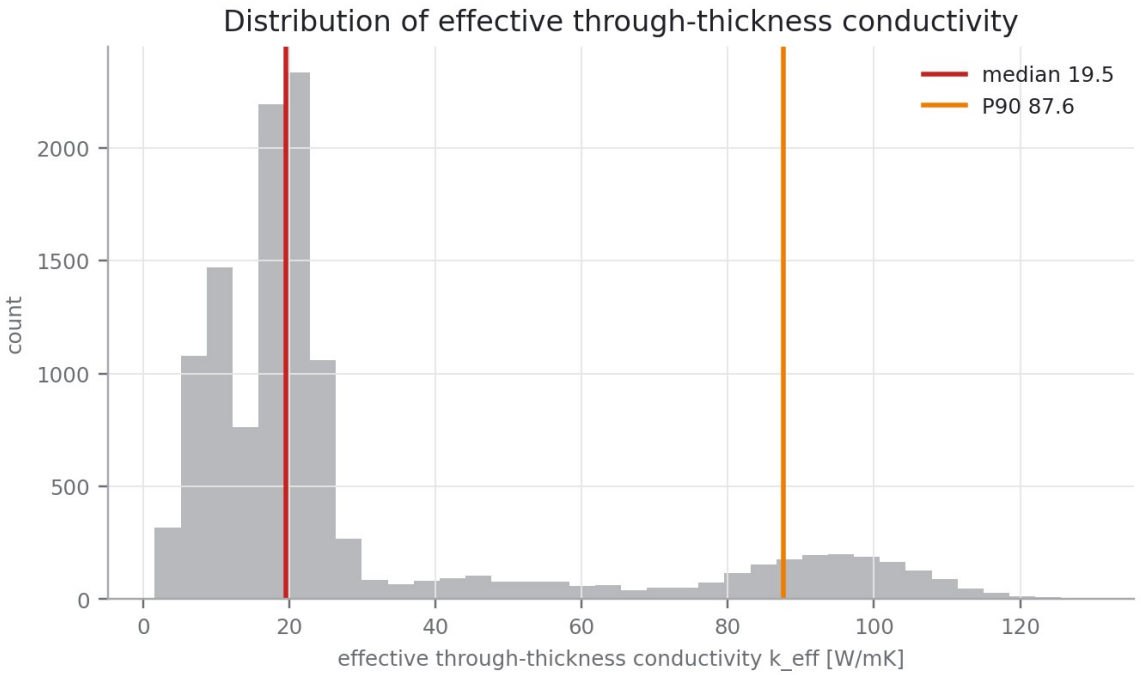


Figure 21: The effective property is strongly shaped by thin resistive terms and not by the best bulk value in the stack.

# 6.3 Distribution of peak overtemperature

Risk is visible in the tail, not only in the median

The median transient peak overtemperature of the modeled population is about 147 K, but the 90th percentile rises to approximately 968 K. The tail matters because thermal damage, coating distress, and local yield often begin in exactly this part of the distribution. The statistical study in this whitepaper is based on 12,000 virtual clad stacks spanning thickness, conductivity, interface, cooling, and pulse windows typical of industrial screening work.

In design review, tail behavior is the reason to keep scatter visible. A materially attractive nominal design may still be operationally weak when its interface and release variability are too wide. The purpose is not to mimic one exact product but to show which variables move the risk distribution and which ones mainly influence the nominal point.

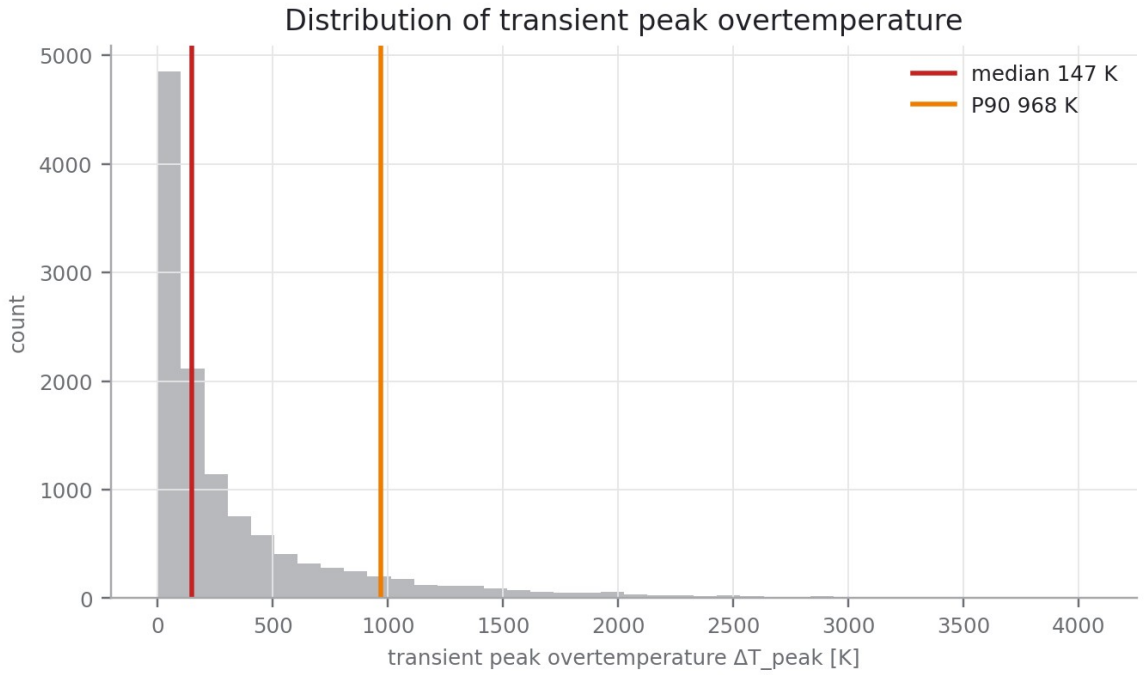


Figure 22: Peak-temperature risk is dominated by the upper tail rather than by the median response.

# 6.4 Distribution of t95

How fast the population equalizes thermally

The homogenization metric t95 shows a strongly skewed distribution. In the modeled study the median lies at about 5.6 s, while the 90th percentile extends to roughly 113.0 s. The statistical study in this whitepaper is based on 12,000 virtual clad stacks spanning thickness, conductivity, interface, cooling, and pulse windows typical of industrial screening work.

This is operationally useful because cyclic systems can now be separated into stacks that reset between pulses and stacks that carry thermal memory into the next cycle. The purpose is not to mimic one exact product but to show which variables move the risk distribution and which ones mainly influence the nominal point.

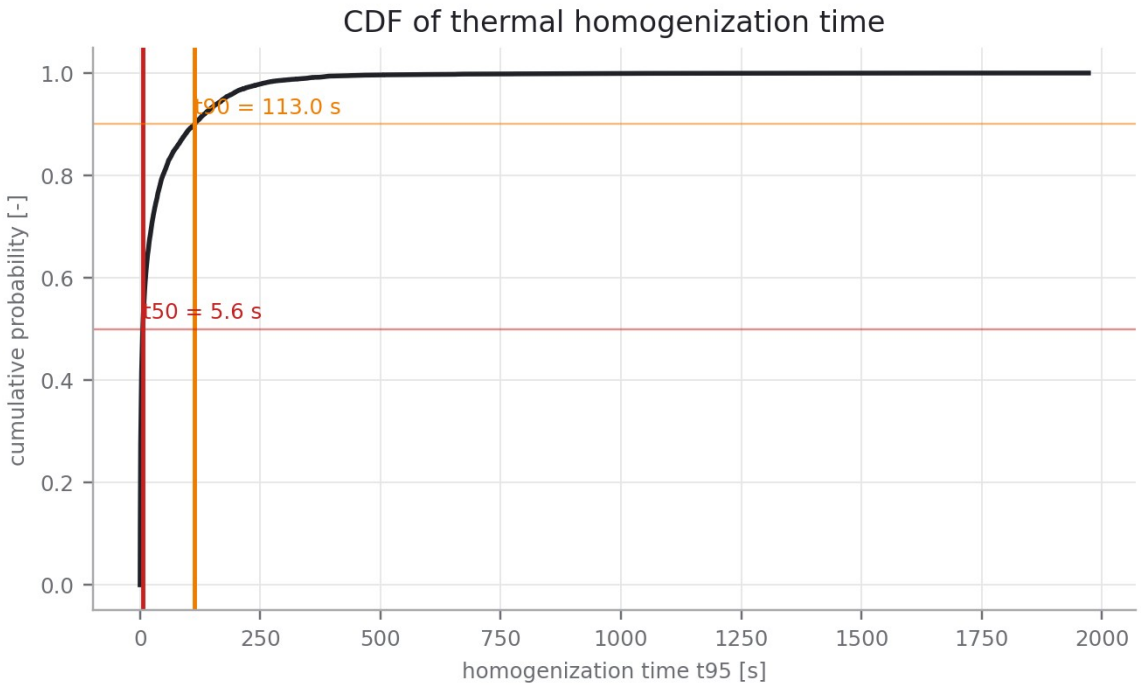


Figure 23: The CDF makes it easy to read typical and tail-side equalization times.

# 6.5 Bond resistance versus peak risk

How interface quality shifts the upper envelope

The population median of ΣR_int is around 0.60 mm²K/W, but the upper tail extends well beyond 2.70 mm²K/W. As the bond resistance rises, the upper envelope of peak risk rises as well. The statistical study in this whitepaper is based on 12,000 virtual clad stacks spanning thickness, conductivity, interface, cooling, and pulse windows typical of industrial screening work.

This plot is especially useful in project reviews because it translates bond quality from an abstract microstructural issue into a directly visible temperature-risk statement. The purpose is not to mimic one exact product but to show which variables move the risk distribution and which ones mainly influence the nominal point.

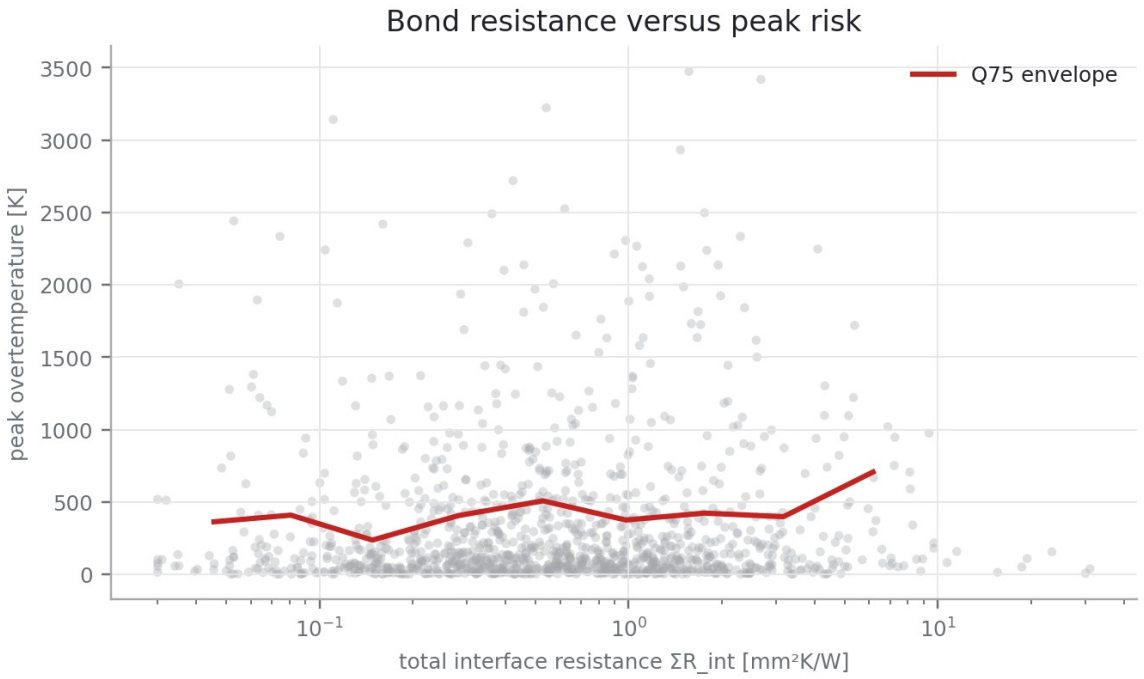


Figure 24: Rising ΣR_int shifts the upper risk envelope even when the bulk stack remains unchanged.

# 6.6 Sensitivity tornado

Which variables move the modeled risk most strongly

Sensitivity ranking in the modeled study shows that cooling coefficient, pulse time, wall thickness, heat flux, and interface resistance together dominate the peak-risk response. The order may change somewhat in a product-specific study, but the principle is robust. The statistical study in this whitepaper is based on 12,000 virtual clad stacks spanning thickness, conductivity, interface, cooling, and pulse windows typical of industrial screening work.

A useful management takeaway follows: invest first where sensitivity is high and controllability is realistic. That usually means boundary-side release, interface evidence, and stack architecture before fine-tuning secondary variables. The purpose is not to mimic one exact product but to show which variables move the risk distribution and which ones mainly influence the nominal point.

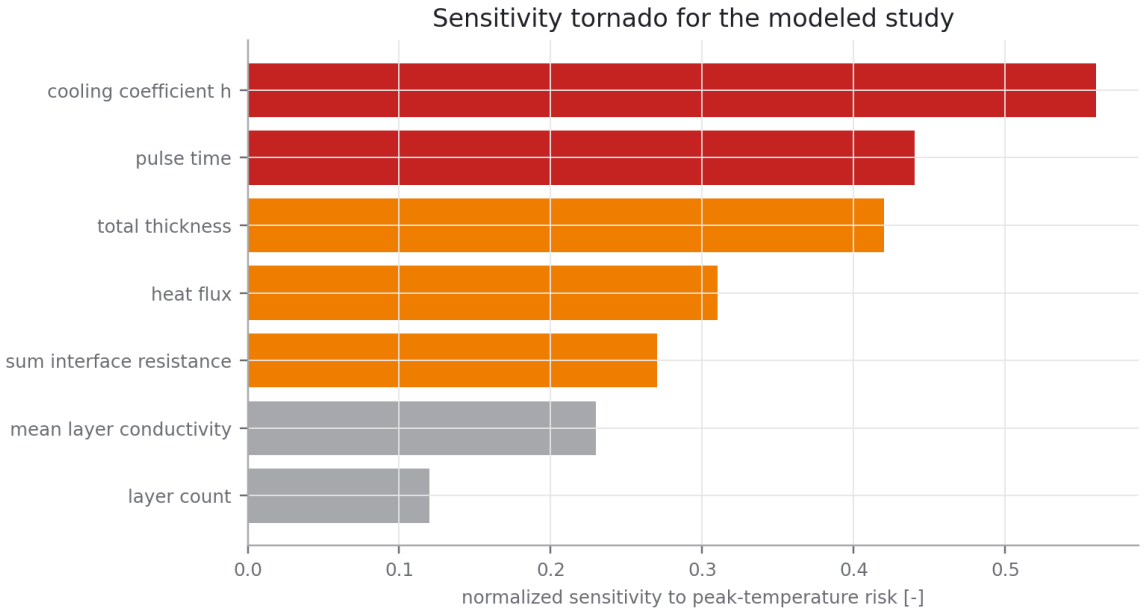


Figure 25: A tornado chart keeps the design conversation focused on the few variables that truly move risk.

# 6.7 Layer count and peak distribution

Why more layers are not automatically better

Adding layers can help when they create a better spreading or functional architecture, but more layers also add more interfaces and more chances for variability. The net effect is therefore conditional rather than monotonic. The statistical study in this whitepaper is based on 12,000 virtual clad stacks spanning thickness, conductivity, interface, cooling, and pulse windows typical of industrial screening work.

The box-plot view is useful because it shows both median shift and scatter change. In practice, layer count is a design decision only when it is read together with interface quality and path logic. The purpose is not to mimic one exact product but to show which variables move the risk distribution and which ones mainly influence the nominal point.

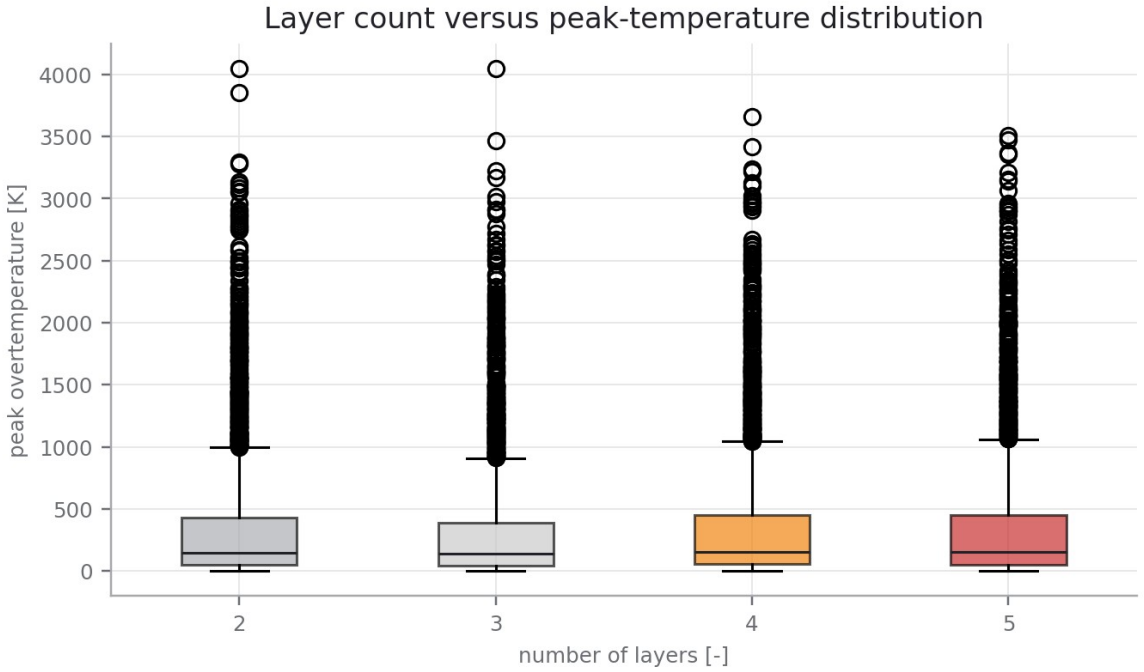


Figure 26: Additional layers can either help or hurt depending on the interface and path architecture.

# 6.8 Correlation structure

Which variables move together in the modeled population

Correlation maps are not causal proof, but they are excellent triage tools. They show which variables tend to shift together across the study and where misleading intuitive interpretations may arise. The statistical study in this whitepaper is based on 12,000 virtual clad stacks spanning thickness, conductivity, interface, cooling, and pulse windows typical of industrial screening work.

For example, a stack with high k_eff may still show high peak temperature when the sink coefficient is poor or when the pulse is too short to benefit from the deeper conductive path. The purpose is not to mimic one exact product but to show which variables move the risk distribution and which ones mainly influence the nominal point.

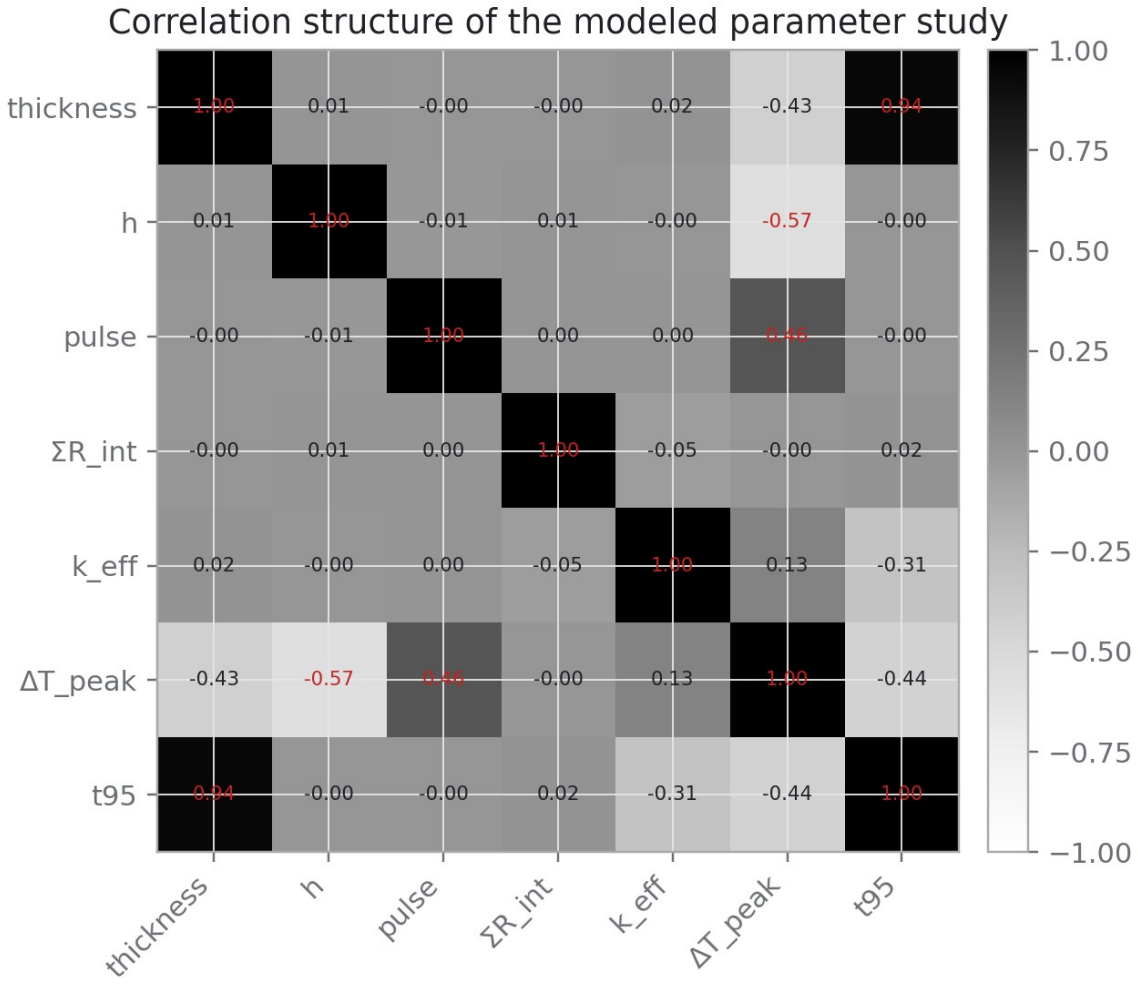


Figure 27: Correlation structure helps prevent simplistic one-variable narratives.

# 6.9 Design window

Combining effective conductivity and cooling-side release

A design window is more useful than a single target value when teams have to balance architecture and cooling effort. The window view shows which combinations of k_eff and h are likely to achieve a smooth thermal release and which combinations are structurally weak. The statistical study in this whitepaper is based on 12,000 virtual clad stacks spanning thickness, conductivity, interface, cooling, and pulse windows typical of industrial screening work.

This kind of map is commercially powerful because it allows design, sales, and customer engineering to discuss trade-offs without pretending that one side of the system can compensate for every weakness on the other side. The purpose is not to mimic one exact product but to show which variables move the risk distribution and which ones mainly influence the nominal point.

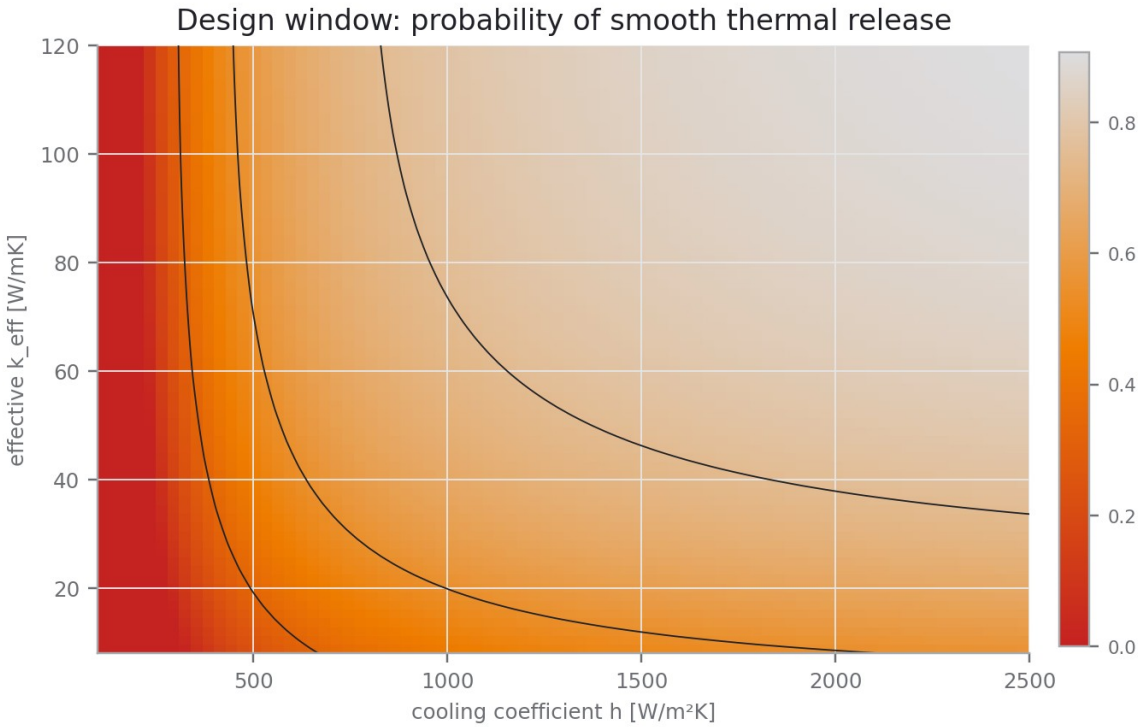


Figure 28: Conductivity and release coefficient must be improved together when the operating point sits in a weak window.

# 6.10 Robustness and uncertainty propagation

From parameter scatter to thermal decision quality

Robustness is the ability of a design to remain acceptable under realistic variation of geometry, bond quality, mounting, and duty. Uncertainty propagation connects this variation directly to temperature and time metrics. The statistical study in this whitepaper is based on 12,000 virtual clad stacks spanning thickness, conductivity, interface, cooling, and pulse windows typical of industrial screening work.

The practical rule is simple: approve the stack against a band, not against a single deterministic point. That band should reflect the variables that are physically relevant and operationally controllable. The purpose is not to mimic one exact product but to show which variables move the risk distribution and which ones mainly influence the nominal point.

- Use distributions for approval where the duty is narrow and the consequence of failure is high.
- Turn uncertain interfaces and release conditions into explicit parameter windows.
- Prefer the architecture with the smaller tail risk, not only the better nominal point.

# 7. Technical renderings and application archetypes

Visualizing the dominant heat path in recurrent product families

The purpose of the renderings in this chapter is operational. They show how the abstract thermal logic of the previous chapters appears in common component families and where the dominant path, local bottleneck, and release condition sit in each case.

This is especially useful in customer communication and concept review. Once the application archetype is visible, arguments about material, process, and test depth become more concrete and easier to align across disciplines.

- Heat spreaders convert local source peaks into wider sink-facing footprints.
- Tubes and rolls are governed by radial area change and contact condition.
- Battery and electronics plates demand both low level and low spread.

# 7.1 Heat spreader and local source load

A classic case of lateral smoothing before release

The heat-spreader archetype is driven by a concentrated source that must be smoothed before the load is allowed to enter the sink leg. A conductive cap layer is useful here because it increases lateral path width and reduces the local source peak. Application archetypes make the thermal path visible and help translate abstract conduction equations into concrete design choices.

The risk arises when the path below the spreader is still too resistive. In that case the hot spot remains trapped above a bottleneck despite the improved lateral field. For customer-facing work this is important because the relevant argument is usually not the material label but the heat-path logic of the part in service.

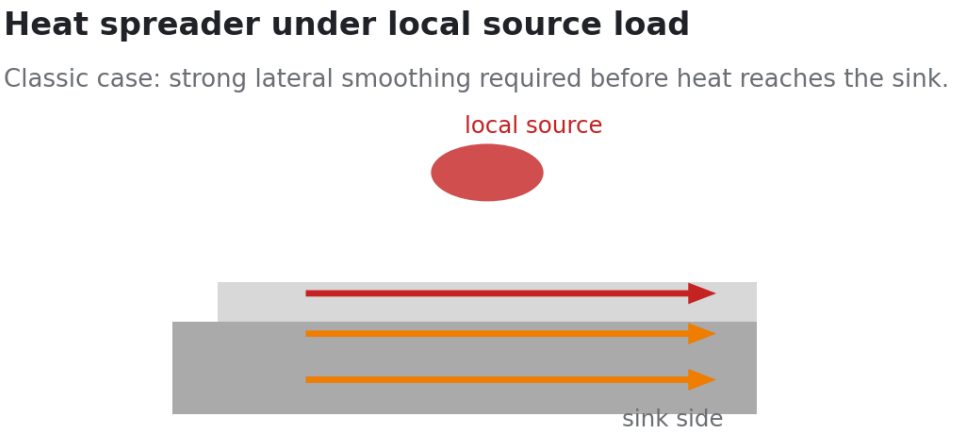


Figure 29: Local source cases benefit from conductive caps only when the vertical release leg remains credible.

# 7.2 Clad tube and radial heat path

When the path wraps around a cylindrical geometry

Clad tubes, liners, and sleeves often release heat radially. The changing radius means changing area, which in turn creates a logarithmic profile and a different weighting of inner and outer layers than in a flat-wall reading. Application archetypes make the thermal path visible and help translate abstract conduction equations into concrete design choices.

For process equipment this matters because thin corrosion-resistant or functional layers may still carry a disproportionate thermal consequence when placed at the inner high-flux side. For customer-facing work this is important because the relevant argument is usually not the material label but the heat-path logic of the part in service.

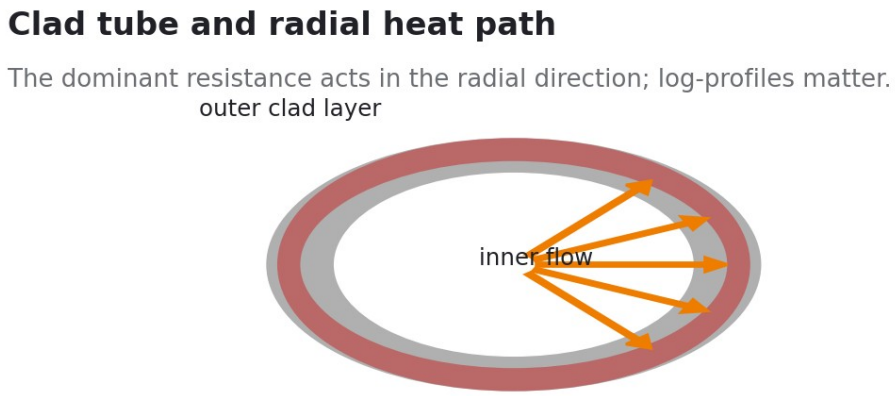


Figure 30: Radial geometry changes the conduction picture and makes cylindrical interpretation mandatory.

# 7.3 Vessel and reactor wall

Distributed process-side load with aggressive service requirements

Reactor and vessel walls often combine distributed thermal duty with demanding chemical or corrosion requirements. The clad concept must therefore satisfy function and temperature at the same time, usually with very different layers carrying those responsibilities. Application archetypes make the thermal path visible and help translate abstract conduction equations into concrete design choices.

The useful engineering question is not whether the functional layer is thermally ideal, but whether its resistance contribution is acceptable within the full release budget of the wall and the real process-side load. For customer-facing work this is important because the relevant argument is usually not the material label but the heat-path logic of the part in service.

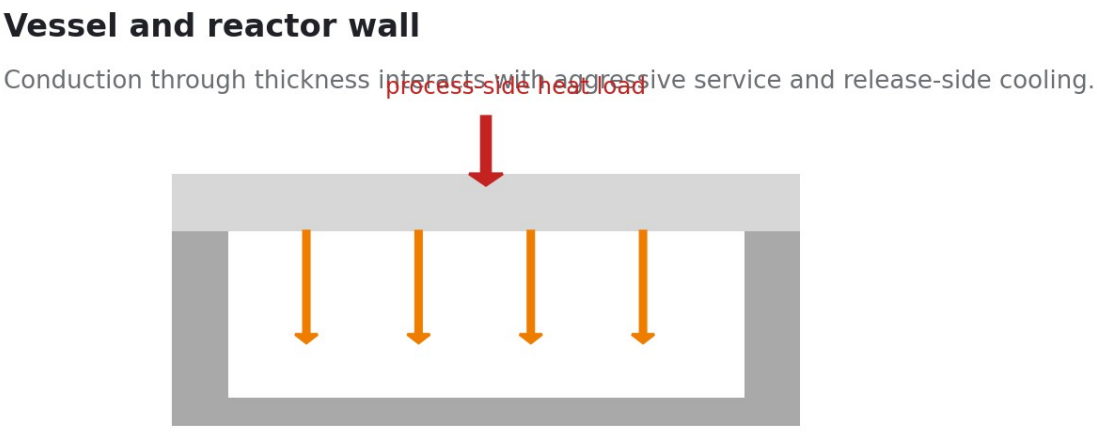


Figure 31: Vessel walls combine process function, wall conduction, and release-side coupling in one path.

# 7.4 Roll or tool under cyclic thermal loading

Pulse stacking, thermal memory, and local contact load

Rolls and thermal tools often see repeated local contact events. The critical issue is whether the system resets between cycles or whether thermal memory accumulates from pulse to pulse. Application archetypes make the thermal path visible and help translate abstract conduction equations into concrete design choices.

The right clad concept is therefore one that controls both the peak during one event and the recovery before the next event. t95 and sink-side coupling become central metrics in this archetype. For customer-facing work this is important because the relevant argument is usually not the material label but the heat-path logic of the part in service.

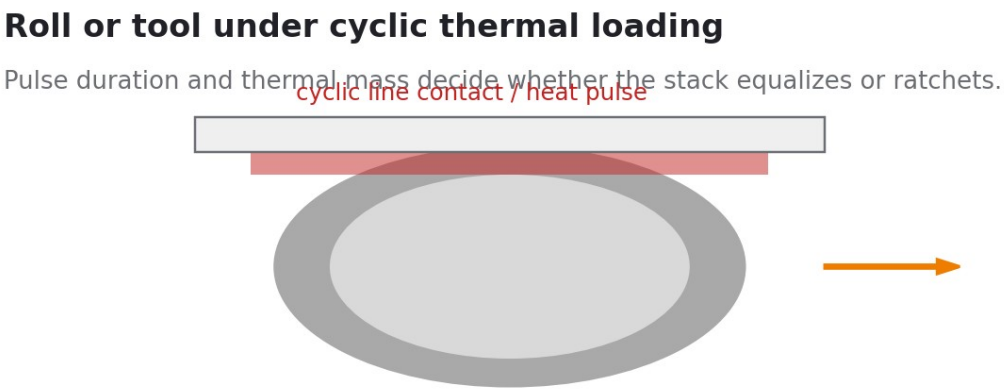


Figure 32: Cyclic tools must be read in terms of both local pulse peak and inter-cycle recovery.

# 7.5 Battery and electronics cooling plate

Low mean temperature is not enough; homogeneity also matters

Battery and electronics cooling plates are archetypes of mixed requirements. They need low temperature level, strong lateral smoothing, and predictable transfer from source array to coolant or base plate. Application archetypes make the thermal path visible and help translate abstract conduction equations into concrete design choices.

In this family the quality metric is often not the absolute maximum alone but also the temperature spread. Clad concepts become attractive when they offer a tunable compromise between spreading, structure, and manufacturability. For customer-facing work this is important because the relevant argument is usually not the material label but the heat-path logic of the part in service.

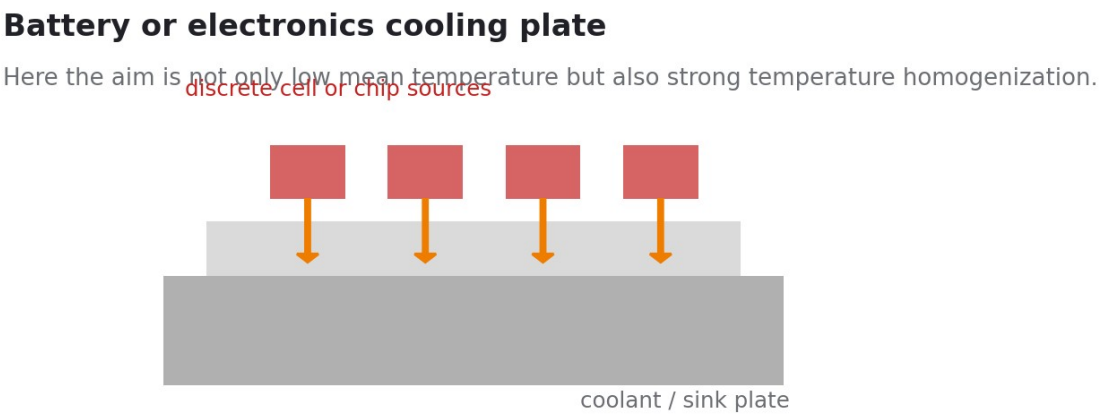


Figure 33: Cooling plates are judged on average level, spread, and repeatability of release.

# 7.6 Typology of boundary conditions

The same stack behaves differently under different release definitions

Application archetypes differ not only by geometry but also by boundary-condition type. Prescribed heat flux, prescribed temperature, finite contact resistance, and convective release create fundamentally different thermal narratives for the same physical stack. Application archetypes make the thermal path visible and help translate abstract conduction equations into concrete design choices.

A well-formed project specification therefore names the load type, the contact description, the sink evidence, and the ambient evidence together. Without that, thermal comparisons between concepts are often misleading. For customer-facing work this is important because the relevant argument is usually not the material label but the heat-path logic of the part in service.

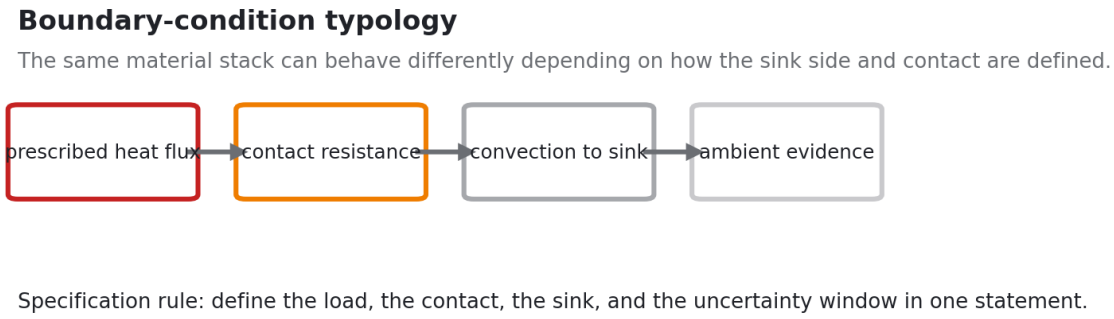


Figure 34: The thermal story changes as soon as the load or sink description changes.

# 7.7 Typical thermal failure patterns

Using field signatures to triage root cause

Common failure signatures include a concentrated hot spot above a local source, a temperature jump at a delaminated interface, edge overheating where spreading distance collapses, and a delayed sink response that indicates insufficient release or excessive thermal mass. Application archetypes make the thermal path visible and help translate abstract conduction equations into concrete design choices.

These signatures are diagnostically valuable because they point to mechanism rather than to symptom alone. They help decide whether the next action belongs in material choice, process control, assembly, or model refinement. For customer-facing work this is important because the relevant argument is usually not the material label but the heat-path logic of the part in service.

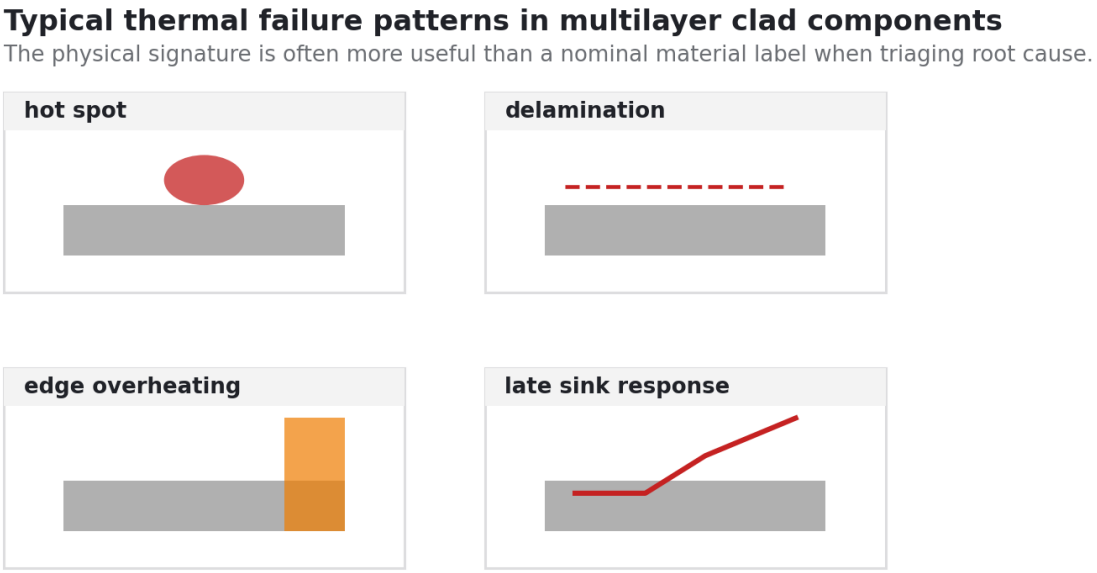


Figure 35: Thermal field signatures are often the fastest route from symptom to likely mechanism.

# 8. Design rules, governance, and QA

Turning thermal logic into an operational release system

A technically strong stack still needs governance. Design rules, procurement language, release criteria, and production monitoring must all refer to the same thermal logic; otherwise the thermal concept is diluted during industrialization.

This chapter translates the preceding physics into a set of practical rules for model depth, layer ordering, cooling-side mounting, scaling, tolerance handling, and specification language.

- Write one thermal statement that design, procurement, and QA can all use.
- Approve against a controlled window, not against one idealized sample.
- Connect field evidence, reduced models, and production metrics.

# 8.1 Decision matrix for model and test depth

Scaling evidence with risk and complexity

Model and test depth should increase with path complexity and release risk. Simple distributed cases can be handled with screening equations and one confirmation test; local-source cases with weak interfaces or strong pulses may require reduced transient models, metrology, and eventually a local 2D or 3D study. The goal of governance is to align design, procurement, testing, and release around one coherent thermal statement.

The important business rule is proportion. Too little evidence creates false confidence, while too much evidence delays viable projects without improving decision quality. Projects become faster and safer when specifications state the path, the interfaces, the sink, the load window, and the admissible drift in one place.

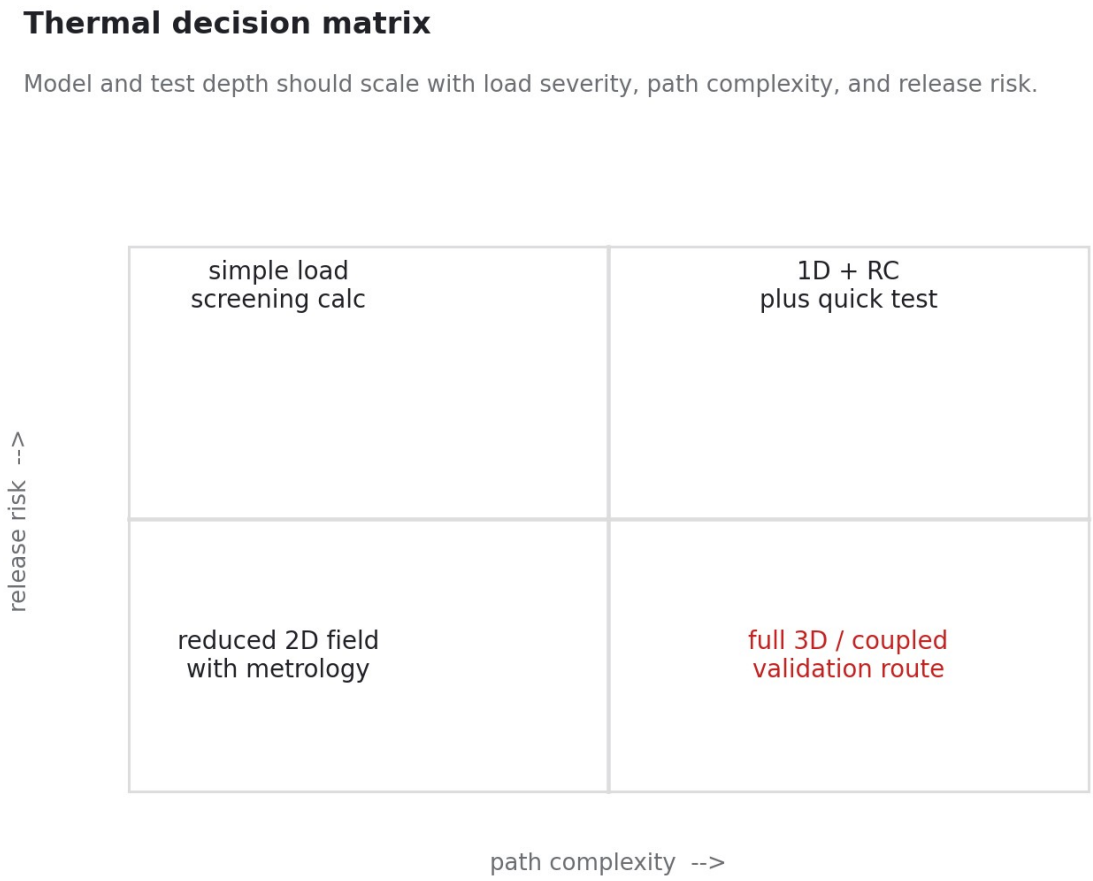


Figure 36: Path complexity and release risk together define how much evidence is proportionate.

# 8.2 Heuristics for layer sequence

Ordering layers according to their thermal role

Layer sequencing is one of the most powerful early levers. Put high-k material where spreading is required, keep thermal barriers as thin as their function allows, and place structurally thick low-k layers where they do not sit directly beneath local heat peaks unless the sink path is very strong. The goal of governance is to align design, procurement, testing, and release around one coherent thermal statement.

Sequence should be argued in thermal-role language: spread, block, carry, protect, or release. That language aligns better with engineering decisions than a simple list of materials. Projects become faster and safer when specifications state the path, the interfaces, the sink, the load window, and the admissible drift in one place.

- Place the spreader near the source if local smoothing is needed.
- Minimize the thickness of thermally expensive functional layers.
- Do not hide a weak sink path under a strong cap layer.

# 8.3 Cooling-side coupling and assembly

Mounting detail as part of the thermal design

The cooling-side connection deserves explicit design ownership. Flatness, torque, clamping concept, contact media, grease or pad selection, and cleanliness all influence 1/h and the effective contact resistance seen by the stack. The goal of governance is to align design, procurement, testing, and release around one coherent thermal statement.

A sound release route therefore couples material approval with assembly evidence. If mounting quality is expected to vary, that variability should be represented explicitly in the thermal acceptance band. Projects become faster and safer when specifications state the path, the interfaces, the sink, the load window, and the admissible drift in one place.

- Treat contact evidence as part of the thermal data package.
- Specify assembly conditions together with thermal acceptance criteria.
- Monitor drift where the release concept relies heavily on contact quality.

# 8.4 Geometric scaling

Why a coupon does not automatically predict a full-size component

Scaling changes path length, spreading distance, edge influence, and sometimes even the effective release coefficient. A result measured on a small coupon therefore cannot be transferred blindly to a large plate, a tube, or a formed component. The goal of governance is to align design, procurement, testing, and release around one coherent thermal statement.

The correct transfer question is not whether the material is the same, but whether the dominant heat path and the relevant dimensionless ratios remain comparable. Projects become faster and safer when specifications state the path, the interfaces, the sink, the load window, and the admissible drift in one place.

- Coupon evidence is strongest when the dominant path remains unchanged.
- Larger geometries often increase spreading distance and reduce apparent uniformity.
- Scale-up should be reviewed against path similarity, not against material identity alone.

# 8.5 Manufacturing tolerances, bond continuity, and drift

Bringing production realism into the release discussion

Thickness tolerance, waviness, local bond discontinuity, and route-specific microstructural drift all change the thermal response band. The effect is usually modest on the nominal point but significant in the tail. The goal of governance is to align design, procurement, testing, and release around one coherent thermal statement.

A robust industrial release therefore identifies which manufacturing variables are allowed to drift, how much they may drift, and which thermal metric is used to detect the drift before customer impact occurs. Projects become faster and safer when specifications state the path, the interfaces, the sink, the load window, and the admissible drift in one place.

- Track the variables that physically move peak temperature and t95.
- Treat bond continuity as a measurable quality feature where possible.
- Translate tolerance into thermal consequence instead of reviewing dimensions in isolation.

# 8.6 QA dashboard and release criteria

A compact metric set for serial control

A useful QA dashboard is compact. It combines one absolute temperature metric, one homogenization or spread metric, one maturity or release status measure, and one drift indicator that can be tracked over time. The goal of governance is to align design, procurement, testing, and release around one coherent thermal statement.

This avoids over-instrumented quality systems and keeps the thermal control loop connected to what really matters in production. Projects become faster and safer when specifications state the path, the interfaces, the sink, the load window, and the admissible drift in one place.

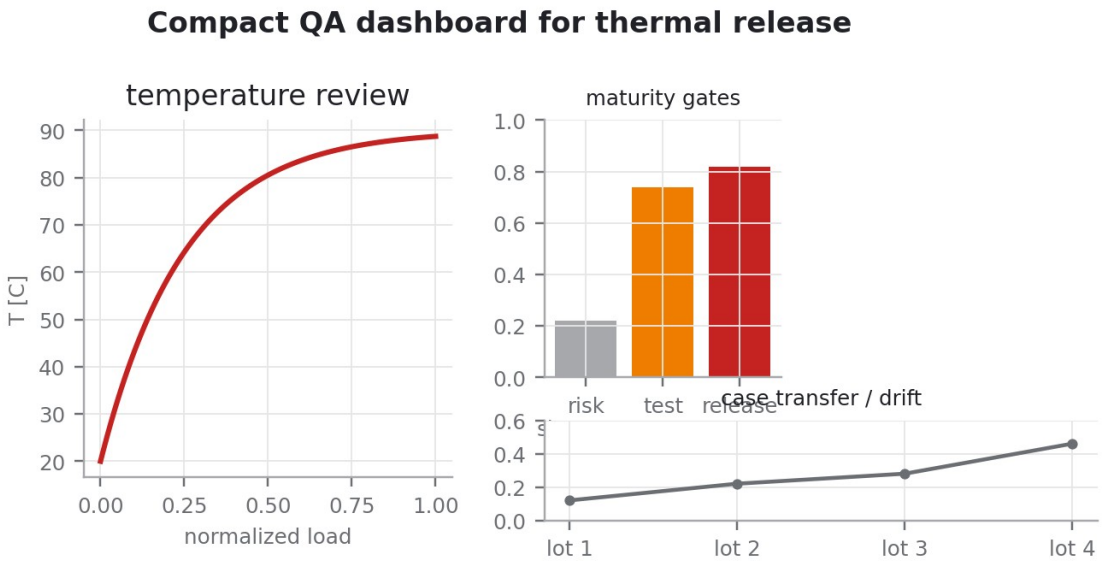


Figure 37: A compact dashboard aligns temperature level, maturity gates, and lot-to-lot drift.

# 8.7 Specification language for procurement and project control

How to write a thermal requirement that survives execution

Procurement and project documents should describe the thermal requirement as a path statement, not as an isolated material callout. That statement should name the load, the stack role of each layer, the interface requirement, the sink-side evidence, the primary metric, and the admissible drift window. The goal of governance is to align design, procurement, testing, and release around one coherent thermal statement.

A specification written in this way is technically clearer and commercially safer because it links supplier obligations to measurable thermal outcomes instead of ambiguous labels. Projects become faster and safer when specifications state the path, the interfaces, the sink, the load window, and the admissible drift in one place.

- Name the primary thermal metric explicitly.
- State the boundary condition and mounting evidence in the same clause.
- Require interface and continuity evidence where the path is bond-sensitive.

# 9. Worked Example A - steady-state plate

A simple plate section under stationary heat flux

Worked Example A illustrates the steady-state logic with a compact plate stack. The objective is not to represent one exact product but to show the calculation path from geometry and material data to resistance budget and temperature estimate.

The example uses a conductive cap, one functional interlayer, and a structural carrier under a uniform source-side heat load. The result is read both numerically and as a temperature profile through thickness.

- Inputs are chosen for engineering plausibility, not for certification use.
- The example highlights the difference between cap benefit and sink-side control.
- The same framework can be reused in quotations and concept studies.

# 9.1 Inputs and formula path

Geometry, material set, and thermal ledger

Input set: cap layer 0.35 mm, functional interlayer 0.45 mm, carrier 1.80 mm, representative conductivities 205 / 45 / 16 W/mK, source-side heat flux 0.12 MW/m², and sink-side cooling coefficient 1400 W/m²K. The total resistance per unit area is formed as the sum of layer conduction and the sink term.

The example shows how quickly the sink-side release can dominate the total budget even when the wall itself is comparatively conductive. That is exactly why assembly and boundary definition belong in early plate sizing.

**Example A ledger**

```
R_cond/A = Σ( t_i / k_i ) = 1.24e-4 m²K/W
R_conv/A = 1 / h = 7.14e-4 m²K/W
ΔT = q" * ( R_cond + R_conv ) ≈ 100 K
```

# 9.2 Result and interpretation

Reading the dominant term in the temperature path

The resulting profile shows a modest drop across the conductive cap, a visible but still controlled jump across the functional layer, and the largest single contribution at the release side. The total rise is acceptable for the assumed operating point, but the example makes the dominant term visible: boundary-side release.

The interpretation is operational. Further improvement should focus on sink evidence and contact quality before costly cap-layer upgrades are considered.

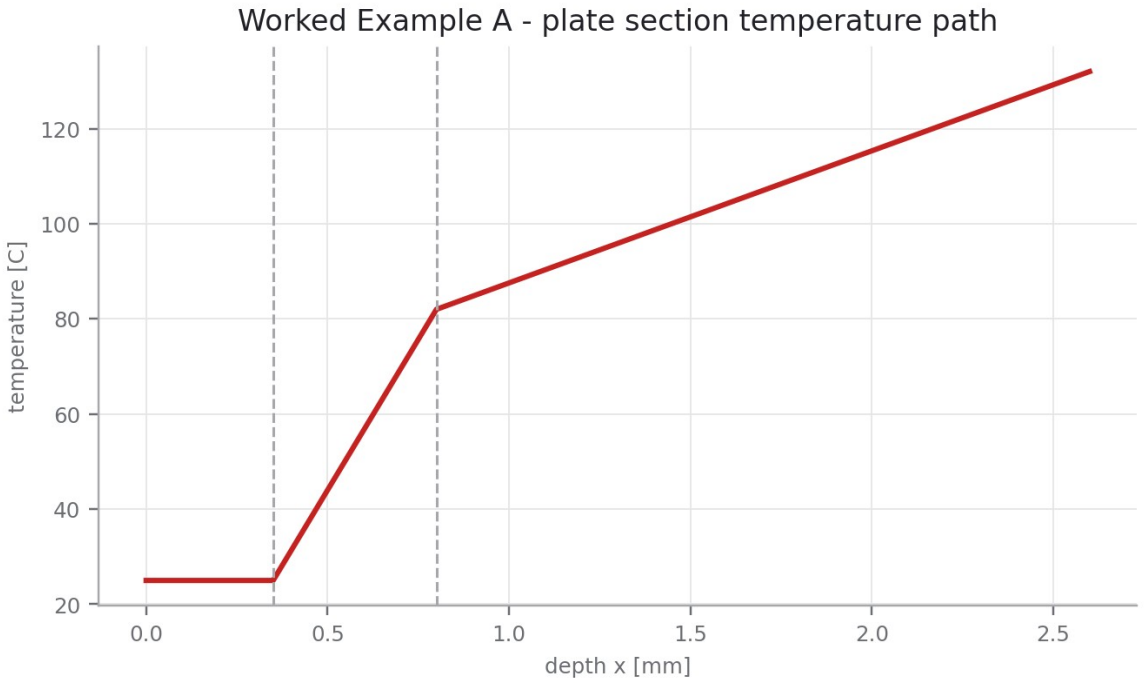


Figure 38: Example A - the profile shows where the thermal head is actually spent.

# 10. Worked Example B - pulse loading

A reduced transient example for a three-layer stack

Worked Example B shifts the focus from resistance to response time. The same general stack logic is now tested under a pulse where the decisive question is whether the wall equalizes during the event and whether the sink-side node remains slow relative to the source-side rise.

A reduced RC description is sufficient for the engineering message and makes it possible to estimate peak temperature, intermediate node behavior, and recovery trend with little computational effort.

- Dynamic storage matters as much as resistance.
- Pulse duration must be compared with the characteristic stack time.
- Reduced models are ideal for concept-side transient sweeps.

# 10.1 Inputs and model setup

Three thermal masses and their connecting resistances

Input set: cap 0.40 mm, interlayer 0.70 mm, carrier 1.60 mm, representative pulse heat flux 0.22 MW/m², sink-side coefficient 700 W/m²K, and a reduced three-node RC model with source-side, mid-stack, and sink-side nodes. The pulse duration is 30 s and the stack starts from a uniform initial temperature.

The model separates rapid source-side rise from slower deeper response. This is enough to decide whether a local peak is short-lived or whether thermal memory will influence the next event.

**Example B dynamic quantities**

```
τ_surface ≈ 5.8 s
τ_mid ≈ 11.5 s
τ_sink ≈ 18.5 s
t95,system ≈ 3 * τ_dominant
```

# 10.2 Time response and interpretation

Peak, lag, and thermal memory

The source-side node rises quickly and approaches its local peak within a few seconds, while the middle and sink-side nodes lag behind. This temporal separation is the hallmark of a pulse-sensitive stack and explains why one sensor location cannot describe the whole event.

For design, the implication is straightforward: if pulse repetition occurs before the sink-side node has substantially recovered, the stack must be reviewed for cumulative heating rather than for isolated events only.

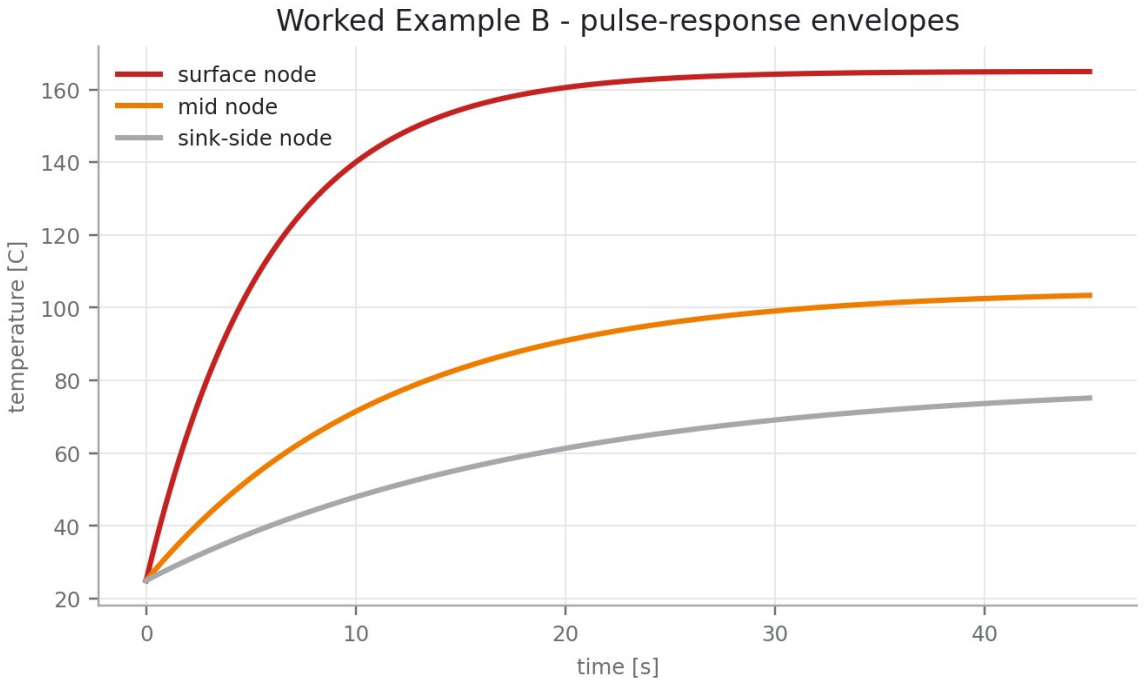


Figure 39: Example B - source-side nodes react quickly, deeper nodes respond with increasing lag.

# 11. Worked Example C - cylindrical wall

A radial conduction example for a clad tube or sleeve

Worked Example C applies the same thermal logic to a cylindrical geometry. The calculation highlights the need to use logarithmic radial resistance instead of a flat-wall shortcut when the wall wraps around a tube or sleeve.

The example is typical for process-side tubes, rolls, or sleeves where the inner and outer surfaces do not carry the same geometric weight in the heat path.

- Area changes with radius and changes the thermal weighting of layers.
- Radial conduction often interacts strongly with contact or fluid-side release.
- A flat-wall shortcut is only acceptable for very small radial thickness relative to radius.

# 11.1 Inputs and formula path

Radial resistance and representative dimensions

Input set: inner radius 12 mm, clad-wall increment 0.8 mm, representative line heat rate 5200 W/m, and a metallic clad layer over a structural base. The radial resistance is calculated with the logarithmic cylinder expression and then combined with the outer release term where required.

The useful message is structural: once the geometry becomes cylindrical, radius placement matters and the thermal consequence of a thin inner or outer layer cannot be judged correctly with a flat-wall analogy.

**Radial relations**

```
R_rad = ln(r_o / r_i) / ( 2π k L )
q' = 2π L k (T_i - T_o) / ln(r_o / r_i)
Temperature varies linearly with ln(r) rather than with r
```

# 11.2 Radial profile and interpretation

What the cylindrical profile tells the reviewer

The radial profile decays logarithmically from the loaded inner side toward the outer side. The slope is steepest where radius is smallest, which is why inner functional layers can have a large influence even when their physical thickness is modest.

Operationally, the example reinforces a governance rule of the whitepaper: match the model to the path geometry first, then discuss material optimization.

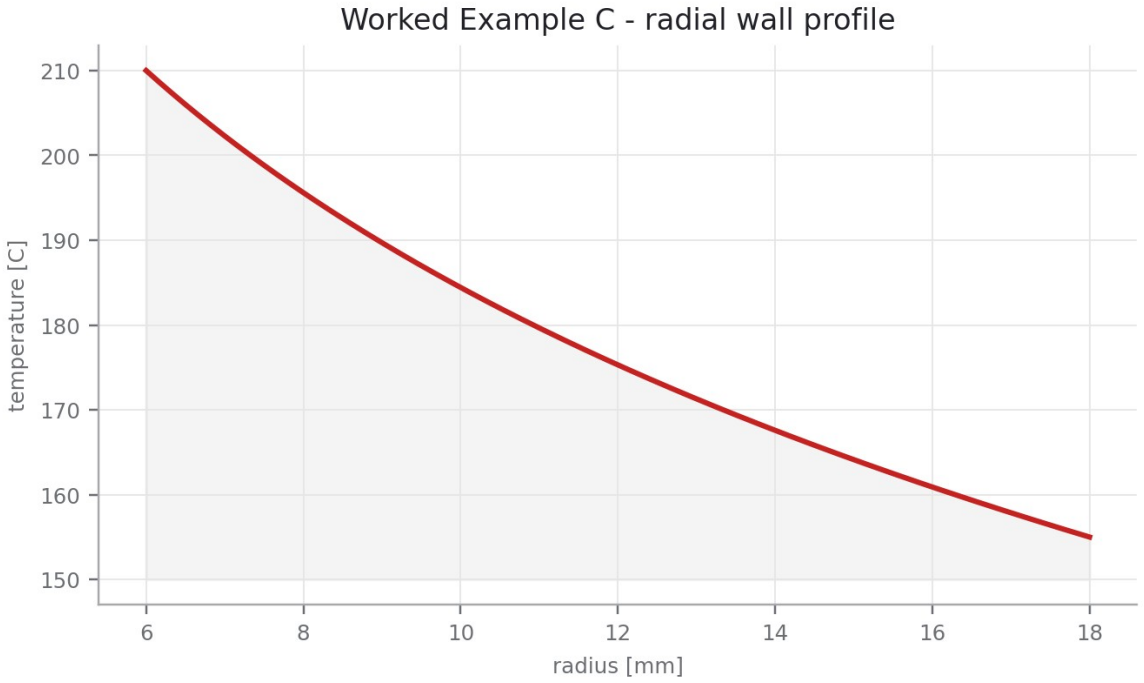


Figure 40: Example C - cylindrical interpretation changes both slope and weighting of the wall sections.

# Appendix A - formula compendium

Compact formula collection for design and review

The appendix condenses the most important equations used in this whitepaper. The intention is not to replace detailed derivations but to provide a compact reference that can be used during concept reviews, quotation work, and internal design discussions.

Equations are grouped by steady-state conduction, interface and convection terms, and transient quantities.

- Use the compact forms for quick review and screening.
- Escalate to field models only when the regime or geometry requires it.
- Keep the governing quantity visible even in more detailed models.

# A.1 Core conduction equations

Steady-state relations for plate-like stacks

These equations cover the most common plate-like or wall-like screening tasks. They are the mathematical core of Chapters 2 and 9 and form the reference point for equivalent-property calculations.

All expressions are written per unit area unless noted otherwise.

**Steady-state conduction**

```
q" = -k_perp dT/dx
R_cond/A = Σ_i ( t_i / k_i )
ΔT = q" * ( R_cond + ΣR_int + 1/h )
k_eff = L_tot / ( Σ_i ( t_i / k_i ) + ΣR_int )
```

# A.2 Interface and convection equations

Contact and sink-side relations

Interfaces and release conditions are often the decisive additions to the simple conduction ledger. These relations are included separately because they are frequently forgotten in early concept work despite their strong influence on thermal level and robustness.

Whenever mounting quality or sink evidence is uncertain, these terms should be treated as variable windows rather than as one nominal constant.

**Interface and release**

```
ΔT_int = q" * R_int
R_conv/A = 1 / h
R_total/A = R_cond/A + ΣR_int + 1/h
Bi = h L / k
```

# A.3 Transient and dimensionless groups

Compact dynamic quantities

The following relations are used to classify pulses, thermal memory, and homogenization behavior. They are particularly useful for Chapter 3 and Worked Example B.

Dimensionless groups do not replace physical units, but they provide a clean way to place a case in regime space.

**Transient quantities**

$$\partial T/\partial t = \alpha \, \partial^2 T/\partial x^2$$

$$\alpha = k / (\rho \, c_p)$$

$$\delta = 2\sqrt{(\alpha \, t)}$$

$$Fo = \alpha \, t / L^2$$

$$\tau \approx R_{th} \, C_{th}$$

$$t95 \approx 3 \, \tau$$

# Appendix B - nomenclature and units

Symbols used in this whitepaper

The nomenclature table collects the symbols used across the steady-state, transient, and statistical chapters. Consistent naming is essential because thermal review often spans management, design, test, and quality functions.

Where appropriate, symbols are given in the compact form used in the formulas above.

| Symbol | Meaning | Unit |
|---|---|---|
| q'' | heat flux density | W/m² |
| q' | line heat rate | W/m |
| k, k_eff | thermal conductivity or equivalent conductivity | W/mK |
| ρ c_p | volumetric heat capacity | J/m³K |
| R_int | interface resistance | m²K/W or mm²K/W |
| h | cooling or convection coefficient | W/m²K |
| L, t_i | characteristic length or layer thickness | m or mm |
| α | thermal diffusivity | m²/s |
| τ | characteristic time constant | s |
| t95 | time to 95% response or homogenization | s |
| Bi | Biot number | - |
| Fo | Fourier number | - |

# Appendix C - representative material windows

Reference conductivity windows for frequent material classes

The table below collects representative conductivity windows used for conceptual classification in this whitepaper. Exact values remain material-, temperature-, and route-dependent and must be replaced by project-specific data where release relevance exists.

The windows are deliberately broad enough to support early option screening.

| Material class | Representative k window [W/mK] | Typical role in clad system |
|---|---|---|
| Copper | 380 - 400 | High-conductivity cap or spreader |
| Aluminium | 190 - 220 | Spreader or lightweight carrier |
| Nickel / Ni alloys | 80 - 95 | Functional layer, corrosion resistance, interface control |
| Austenitic stainless steel | 14 - 18 | Structural or corrosion-resistant carrier |
| Titanium | 6 - 8 | Lightweight structural or corrosion-resistant layer |
| Ceramic interlayer | 2.5 - 3.5 | Barrier or functional insert |
| Polymer film | 0.2 - 0.6 | Electrical isolation or specialist function with thermal penalty |

# Appendix D - project checklist and review gates

A compact review sheet for concept and release meetings

The checklist is intended for project reviews, supplier discussions, and release boards. It converts the whitepaper logic into a short operational gate sequence.

A project does not need every item at concept stage, but the sequence helps prevent blind spots.

| Gate | Key question | Expected evidence |
|---|---|---|
| G1 | Is the dominant heat path identified? | Stack rendering or path sketch |
| G2 | Are the boundary conditions specified? | Load, contact, sink, and ambient statement |
| G3 | Is the primary thermal metric fixed? | Peak, spread, t95, or drift criterion |
| G4 | Is interface quality treated explicitly? | Bond route, continuity, or parameter window |
| G5 | Is the chosen model proportional to the risk? | Screening / RC / 2D / 3D rationale |
| G6 | Is calibration or validation planned? | Sensor concept and validation window |
| G7 | Are tolerances linked to thermal consequence? | Drift logic and acceptance band |

# Conclusion and strategic recommendations

What the thermal reading means for engineering and business

Heat distribution in multilayer clad materials is an architectural problem. The answer never sits in a single conductivity value; it sits in the ordered path made of layers, interfaces, release conditions, and time scale. The strongest projects are the ones that classify this path early, choose the simplest sufficient model, and tie approval to a small set of meaningful thermal metrics.

For Revolit the strategic implication is clear: position clad systems not only by material pairing but by controlled thermal architecture. In commercial terms, this means selling a validated path statement rather than a raw nominal value. In engineering terms, it means using reduced models, renderings, and statistical windows together so that robust decisions are made before expensive late-stage corrections become necessary.

- Make path identification the first step of every technical review.
- Treat interface quality and sink evidence as first-class design variables.
- Prefer robust windows over nominal one-point arguments.
- Use reduced models and statistics to shorten commercial and engineering cycles.

# Contact and document note

Revolit GmbH

Author and publisher: Revolit GmbH, Samstagernstrasse 41, 8832 Wollerau, Switzerland. This whitepaper has been prepared as a professional working document for technical positioning, internal coordination, customer communication, and strategic project structuring. The contents may be used in projects as a methodological reference framework.

For project-specific deepening, Revolit should extend the whitepaper with temperature-dependent material data, real bond characterization, geometry-true boundary conditions, and - where local hot spots dominate - a validating 2D or 3D field model.

**Document note**

This edition combines fundamentals, modeled statistics, renderings, and worked examples into one integrated sizing framework. Project-specific decisions should be supplemented with real boundary conditions and measurement data.