N E W S L E T T E R

Why an Aircraft Engine Is Not One Asset — It Is a Portfolio of Modules

June 2026 | Shannon Ackert

Most investors talk about “the engine” as though it were one asset. That shortcut is useful, but economically incomplete.

The market buys, leases, finances, and appraises a commercial aircraft engine as a single unit. But in maintenance terms, it behaves more like a portfolio of modules. Each module has its own deterioration pattern, LLP position, restoration threshold, and cost profile.

That distinction matters because whole-engine measures can conceal the real exposure. Two engines of the same type, age, thrust rating, and time since shop visit can carry very different forward maintenance burdens depending on which modules the shop has restored, which LLPs remain exposed, and which section is likely to drive the next shop visit.

That modular lens is more important in today's operating environment. With engine shop capacity tight, spare engines expensive, new aircraft deliveries delayed, and some new-generation engines facing durability issues, operators are increasingly focused on targeted module work, quick-turn repairs, and green-time solutions to preserve aircraft availability. The economic question is no longer simply whether an engine is serviceable. It is the amount of useful operating time each module can still support, and the associated forward cost.

KEY TAKEAWAYS

- An engine trades as one asset, but its maintenance economics live in the modules. Fan, compressor, combustor, turbine, gearbox, and LLP positions do not deteriorate or reset on the same clock.
- A shop visit rarely resets the whole engine. It restores selected modules at different depths depending on condition, LLP life, performance margin, and commercial objective.
- “Fresh from shop” is not a complete economic description. The relevant questions are which modules the shop restored, what LLP life remains, and which limitation will bind first.
- Green-time modules prove the market already thinks modularly. Operators use serviceable modules to buy defined future utility without fully restoring the whole engine.
- Lessors should test reserve adequacy by module exposure, not just by headline $/FH rate.
- A blended reserve rate may be convenient, but it can mask where the next major cash event will actually occur.

Architecture Defines the Module Map

Turbofan architecture varies by engine. Many common turbofan engines use a two-spool architecture, separating the low-pressure and high-pressure rotating systems. A three-shaft engine adds an intermediate-pressure system. A geared turbofan adds a reduction gearbox between the fan and the low-pressure drive system.

As **Figure 1** illustrates, these architectures organize the engine's rotating systems differently. Those differences matter because they change the module map. They can change which assemblies planners track, how operators and MROs organize repair exposure, and where future cash obligations sit.

But architecture does not change the central economic logic. Operators and MROs do not treat the engine as a single homogeneous asset. They maintain modules, assemblies, and LLP positions that deteriorate at different rates, require different restoration depths, and create different future cash obligations.

In other words, architecture defines the module map. The module map defines the maintenance exposure.

Figure 1: *Common Turbofan Architectures and Their Modular Building Blocks*

Two-spool, three-shaft, and geared turbofan architectures organize their rotating systems differently. The broader economic point remains the same: the market may trade an engine as one asset, but its maintenance economics reside in the modules.

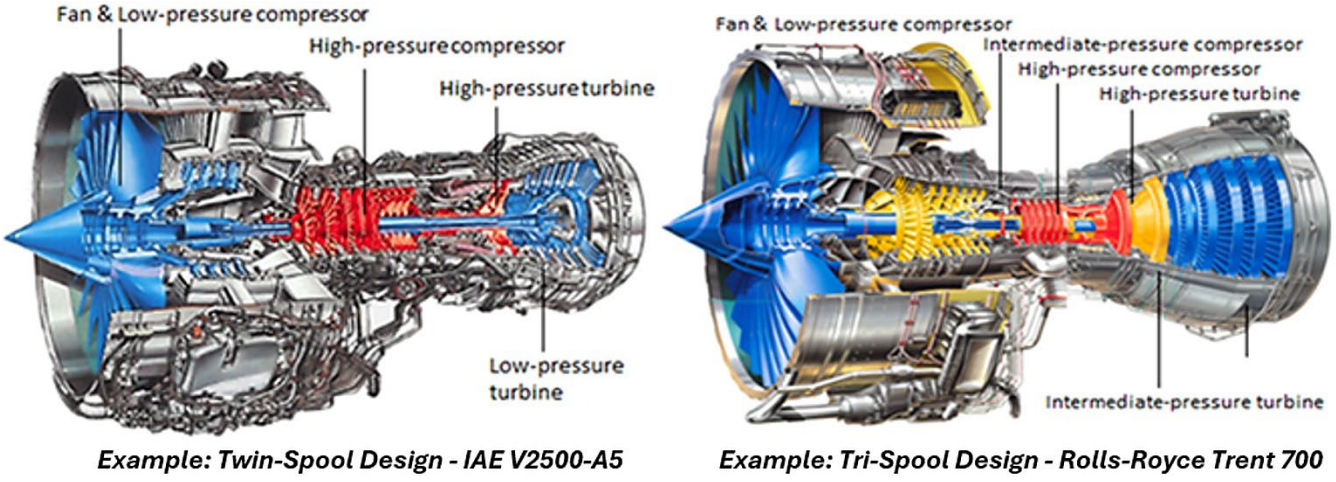

Source: *Representative two-spool and three-spool turbofan module architecture. Adapted from Rolls-Royce, The Jet Engine, 5th ed., John Wiley & Sons, 2015.*

The Engine as a Portfolio of Modules

Figure 2 breaks a generic twin-spool engine into major sections: Fan/LPC, HPC, combustor, HPT, LPT, and, where applicable, accessory gearbox or drive-system assemblies. Maintenance programs can further divide each major section into modules, subassemblies, and life-limited parts.

Figure 2: *Generic Twin-Spool Engine Module Architecture*

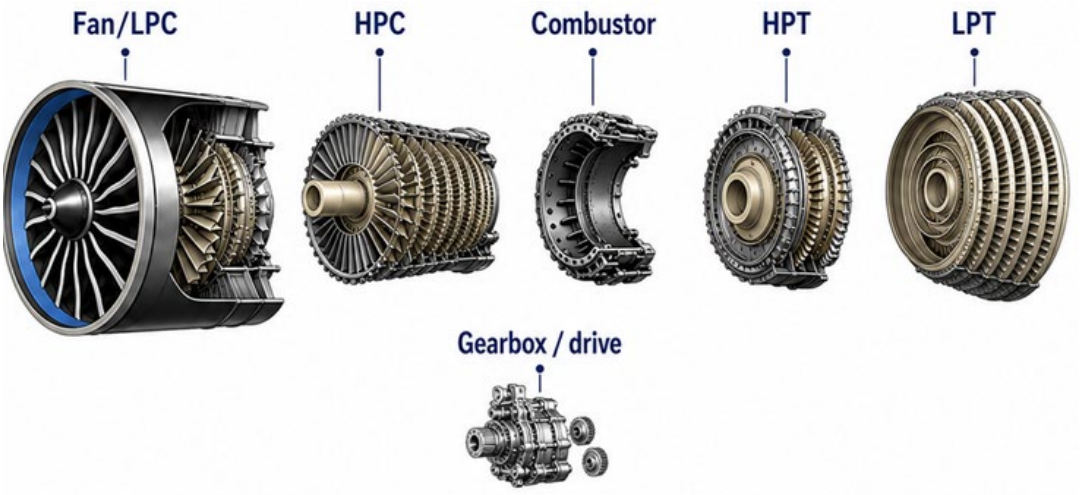

Source Author's Illustration: Major turbofan engine module groupings used for workscope analysis.

Those modules do not deteriorate at the same rate, enter the shop for the same reasons, or consume maintenance capital on the same cycle. Some sections remain relatively stable for long periods. Others, particularly the hotter and more stressed portions of the core, tend to drive earlier and more expensive restoration work.

Future maintenance exposure is therefore not one number. It is a sequence of module-specific cash events, each with its own timing, severity, LLP position, and restoration requirement. That is why analysts should view engine value, lease economics, maintenance reserves, and redelivery exposure through a modular lens rather than through whole-engine time since shop visit alone.

Cold Section, Hot Section, and the Economic Clock

Figure 3 provides a second way to understand an engine's economic behavior: by looking beyond individual modules and separating the architecture into its cold-section and hot-section environments. This distinction matters because the engine does not deteriorate uniformly. The modules exposed primarily to airflow and compression face a different operating environment than the modules exposed to combustion gas, turbine temperatures, and higher thermal stress.

Figure 3: *Engine Architecture Viewed Through the Cold-Section and Hot-Section Lens*

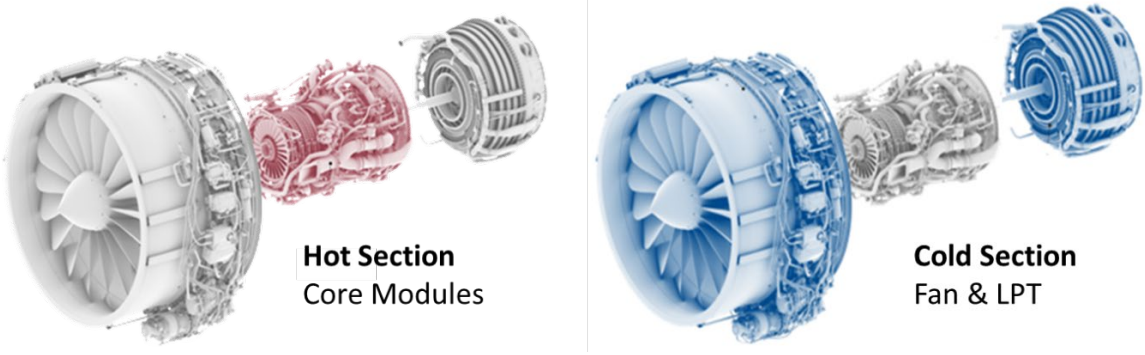

Source: CFMI (LEAP-1A)

The cold section generally includes the front end of the engine: the fan and compression system that moves and pressurizes air before combustion. The hot section begins with combustion and extends through the turbine system that extracts energy from the high-temperature gas stream.

That distinction matters because the hot section experiences the most severe thermal and material stresses and is closely tied to EGT-margin erosion and recovery. Work in the fan or other colder sections can still be necessary for serviceability, durability, and overall performance. But when the operator intends a shop visit to recover EGT-margin and extend on-wing life, the shop often concentrates the most consequential work in the hotter core modules, particularly the combustor and high-pressure turbine.

Thermal deterioration alone, however, does not drive timing. LLP life, hardware condition, records, configuration, and module-specific defects can also drive the next event. The hot section often deserves the closest economic attention. Analysts should not treat it as the only exposure.

Shop Visits Are Modular by Design

A modular lens makes the shop visit easier to interpret. A workscope rarely restores the engine to full uniformity. It addresses selected modules at selected depths.

Some events are limited-scope repairs or quick-turn interventions that correct a specific defect, replace a selected module, or preserve serviceability until a larger workscope makes economic sense. Other events use deeper performance restorations or overhaul-level workscopes to buy a longer interval.

Commercially, the difference is critical. A limited intervention may return the engine to service, but it may not materially change the next binding event. It buys time. It does not necessarily reset the engine.

Figure 4 makes that point clearly. Across the three illustrated shop visits, the shop does not restore the Fan/LPC, HPC, Combustor/HPT, and LPT evenly. The work shifts by module and by event because module conditions, LLP positions, performance needs, and commercial objectives do not align perfectly across the whole engine.

Figure 4*: Illustrative Progression of Module Workscope Levels Across Successive Shop Visits*

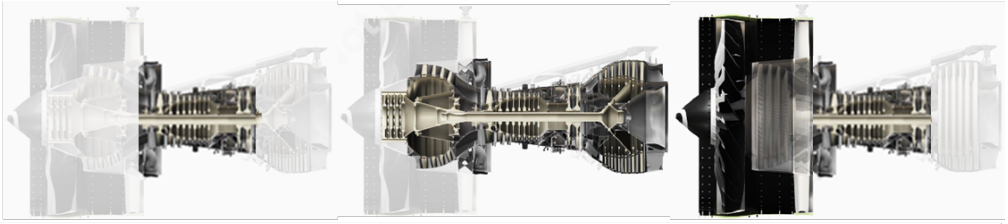

Major Module	SV1 Workscope Level	SV2 Workscope Level	SV3 Workscope Level
Fan / LPC	Minimum Level	Minimum Level	Performance Level
HPC	Performance Level	Overhaul Level	Performance Level
Combustor/HPT	Performance Level	Overhaul Level	Performance Level
LPT	Minimum Level	Performance Level	Minimum Level

Source: Author; illustrative workscope allocation. Not engine-type specific.

That pattern is not random. It reflects how engine economics actually work. Not all modules deteriorate at the same rate, not all modules matter equally in every shop visit, and not all modules reach their LLP limits at the same time.

Minimum, Performance, and Overhaul Workscopes

Minimum level. A minimum-level workscope preserves serviceability, corrects specific defects, and returns the module to operation at the lowest necessary cash outlay. Economically, it is tactical spend: enough to keep the engine operating, but not enough to materially reset that module's future maintenance burden.

Performance level. A performance-level workscope goes further. The primary objective is often to recover EGT margin, restore operating performance, improve reliability, and buy a useful next interval.

EGT margin is the engine's thermal headroom: the difference between the certified temperature limit and the temperature required to produce a given thrust rating. As compressor efficiency, sealing, clearances, combustor condition, and turbine hardware deteriorate, the engine must run hotter to do the same work. That shrinking margin often determines how much useful on-wing life remains.

Overhaul level. An overhaul-level workscope is the most comprehensive and capital-intensive form of intervention. The shop restores the module through broader disassembly, inspection, repair, replacement, and requalification. EGT-margin recovery may remain an important objective, particularly in the hotter core modules, but the broader purpose is to restore module condition, durability, reliability, and future utility. Economically, this is the closest thing to a module-level reset.

The key issue is not whether the engine went to the shop. The better question is whether the modules that matter most to future utility received sufficient restoration to change the engine's forward cash profile.

Performance Restoration and LLP Life Are Different Economic Clocks

Performance restoration and LLP replacement are different economic events. One restores thermal or mechanical capability. The other restores certified life. A shop visit may do one, both, or only part of each.

That distinction matters because an engine can leave the shop with improved EGT margin but still carry meaningful LLP consumption. Conversely, LLP replacement may extend certified life without fully resolving the performance deterioration that triggers the next removal. In practice, the economic objective is often to align the performance restoration interval with the remaining LLP life — or LLP stub-life — and the engine's intended operating horizon. The engine's economic condition, therefore, depends on both clocks: the performance clock and the LLP clock.

For financiers and lessors, time since shop visit is only the starting point. The better questions are which modules the shop restored, how much performance margin it recovered, which LLPs the shop replaced or left in place, and which limitation is likely to bind first at the next event.

Why "Fresh from Shop" Does Not Mean Reset to New

A modular lens makes "fresh from shop" an incomplete description. The phrase tells you an event occurred. It does not tell you what the shop restored, what it left untouched, what LLP life remains, or which module is likely to drive the next cash event.

Figure 5*: Why "Fresh from Shop" Does Not Mean Fully Reset*

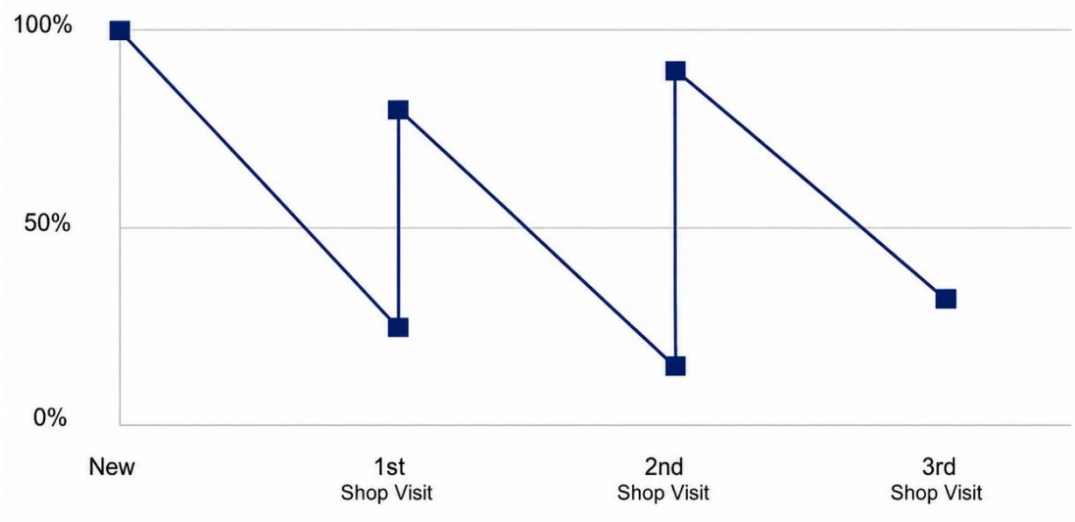

Source: Author; illustrative composite maintenance-utility index.

Figure 5 illustrates a composite maintenance-utility pattern. From new, the engine's maintenance status declines as performance margin erodes, parts age, and service consumes life. After each shop visit, the status improves, but it does not return to its original state. The reason is straightforward: shop visits restore selected portions of the engine, not the entire asset.

A shop visit may restore the HPC and the combustor/HPT to performance or overhaul level, while the Fan/LPC, LPT, or accessory sections receive only minimal work. Unless the shop replaces the entire LLP stack, the engine also carries embedded life consumption from parts left in place. In other words, a shop visit may restore performance capability without fully replenishing the engine's LLP position.

The better question is not whether the engine is fresh from shop. It is fresh in which modules, to what standard, and for how long?

Green-Time Modules: The Aftermarket Proof Point

Green-time modules are one of the clearest real-world expressions of modular engine economics. Operators and MROs do not treat the engine as a single indivisible asset. They assess it module by module, evaluating each section for remaining life, repair cost, configuration, LLP position, records, and expected future utility.

Figure 6: Illustrative Engine Build-Up from Inventory of Green-Time Modules

MROs can use stocked serviceable modules to replace timed-out or uneconomic modules during a shop visit. In this example, the combustor and high-pressure turbine modules no longer support the intended operating interval, so the shop replaces them with green-time modules from inventory. The engine exits the shop with the required remaining utility, without requiring a full restoration of every module.

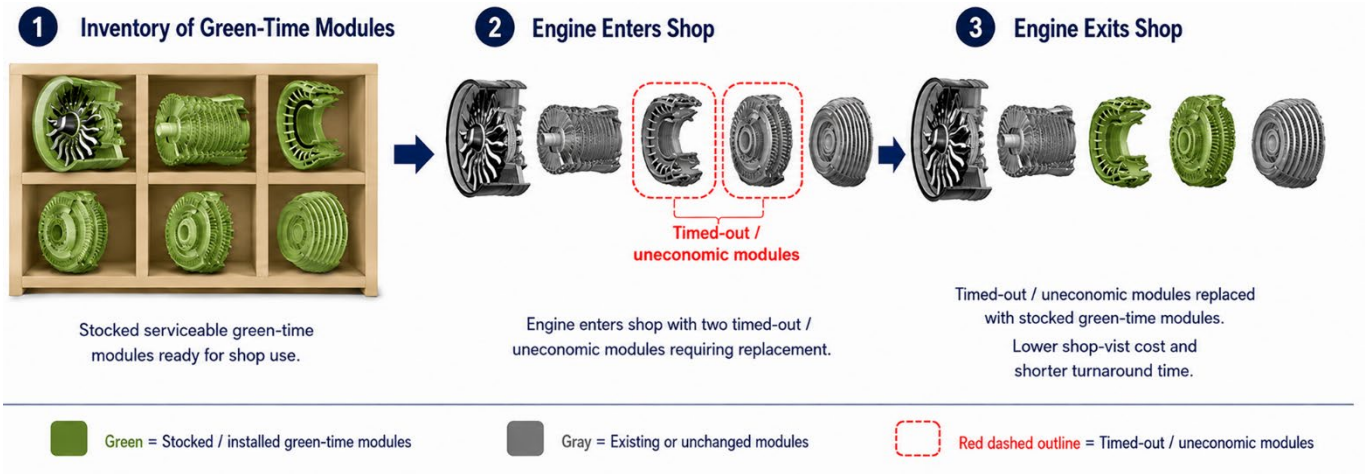

Source: Author

The value of green-time modules increases as an engine program matures. Aircraft retirements, lease transitions, spare-engine rationalization, and teardowns create a larger pool of serviceable used material. In that environment, the operator may not need a fully restored engine. It may need an engine capable of supporting a defined commercial event.

That commercial event may be a lease return, fleet retirement, aircraft transition, spare-engine requirement, replacement-aircraft delay, or planned follow-on shop visit. In those cases, the question is not whether the engine has maximum life. It is whether it has enough reliable utility for the mission ahead.

Green-time modules let the MRO or airline target the shop visit. Rather than replacing or restoring every major life driver, the MRO or airline can focus investment on the modules that no longer support the intended operating plan. In **Figure 6**, the fan/LPC, HPC, and LPT remain unchanged, while the MRO replaces the combustor and HPT with stocked serviceable modules.

Using green-time modules is not simply a shortcut; it requires disciplined technical and commercial review. Before installation, MROs and operators must verify module compatibility, documentation, alignment with the engine build standard, LLP life, and expected time-on-wing. Poorly matched green-time material can create future maintenance exposure, lease-return issues, or residual-value weakness.

When operators use green-time modules correctly, they are not merely a parts-cost tactic. They are a lifecycle-management tool. They are most compelling where two conditions intersect: the operator needs a defined remaining interval, and the aftermarket has enough serviceable used material to support that interval economically.

The broader lesson is simple: a good shop visit does not always buy maximum life. It should buy the right amount of life for the mission ahead.

Why Engine Reserves Need a Module-Level Test

If engines deteriorate by module, if MROs execute shop visits by module, and if maintenance status recovers only partially by module, then lessors should also test engine reserves by module.

A lease may quote a reserve rate as a single dollar amount per flight hour, but the underlying maintenance burden does not build evenly across the engine. The hotter core modules — the high-pressure compressor, combustor, and high-pressure turbine — typically drive the most consequential restoration work. LLP consumption is not uniform across modules, and individual sections move toward their next economic event on different clocks.

The lease may collect the reserve as a single amount, but economically, it funds a sequence of distinct future cash obligations. **Table 1** illustrates the point using a generic $200-per-flight-hour blended engine reserve rate. The purpose is not to prescribe a universal allocation. It shows how a single all-in reserve figure can mask very different module-level exposures.

Table 1: *Illustrative Module-Level Allocation of an Engine Maintenance Reserve Rate*

Module bucket	Allocation / $ per FH	Representative restoration interval	Economic purpose
Fan / LPC	8% / $16 per FH	60,000 FH	Longer-cycle cold-section exposure; support for serviceability and durability.
High-Pressure Compressor (HPC)	12% / $24 per FH	20,000 FH	Compressor efficiency, blade/vane condition, erosion, fouling, and restoration of core airflow performance.
Combustor	10% / $20 per FH	20,000 FH	Hot-section durability exposure: liner, dome, fuel nozzle, and thermal distress management.
High-Pressure Turbine (HPT)	48% / $96 per FH	20,000 FH	Primary hot-section cost driver: EGT-margin recovery, blade/nozzle durability, cooling effectiveness, and life consumption.
LPT	20% / $40 per FH	40,000 FH	Turbine-section restoration exposure and module-specific life consumption.
Accessory gearbox/drive assemblies	2% / $4 per FH	40,000 FH	Drive-system and accessory exposure, where applicable.

Source: Author; illustrative allocation based on a generic blended reserve rate of $200 per flight hour. Not engine-type specific. Restoration intervals are illustrative; the HPC, combustor, and HPT are shown on a common 20,000 FH cycle to align with the worked example in the text, though in practice each module follows its own deterioration and LLP-driven schedule.

The table is not a mere accounting split. It simplifies how future maintenance exposure may sit across the engine. It also explains why a headline dollar-per-flight-hour rate alone cannot prove reserve adequacy.

Two engines may have the same all-in reserve rate yet carry very different forward exposure if their module condition, LLP status, and likely next workscope differ. What matters is not only how much reserve the lease collects in total, but whether the reserve is sufficient for the modules most likely to drive the next major cash event.

Lessor risk arises when the lease collects reserves at a single undifferentiated engine rate but allows reimbursement under a broad "engine performance restoration" trigger. Maintenance reserves are not only a pricing mechanism; they are also a protection mechanism. If reimbursement is not tied to the modules actually restored, the lessor may release the reserve funds needed to cover future exposure in the unrestored modules.

Assume, for example, an engine reserve rate of $200 per flight hour and 20,000 flight hours of utilization. The reserve account would accrue $4.0 million. If the lessee performs a qualifying core performance restoration and the lease allows reimbursement from the full accrued engine reserve fund, the lessee may seek to draw down as much as the full $4.0 million.

Economically, however, the lease did not accrue the full $4.0 million solely for the restored core modules. Based on the illustrative allocation in **Table 1**, the HPC, combustor, and HPT together account for 70% of the reserve rate, or $140 per flight hour. Over 20,000 flight hours, which equate to $2.8 million. The

remaining $1.2 million economically supported the Fan/LPC, LPT, accessory gearbox, and other module-level exposures that the shop may not have restored during the shop visit.

If the lease reimburses the full reserve balance for a partial or module-specific restoration, it may leave the lessor with an engine that has consumed life across multiple modules but with no remaining reserve balance to support those future obligations. In effect, a narrower workscope may consume funds that have been accrued for one maintenance exposure.

A module-level reserve framework helps reduce this exposure. It links reimbursement to the portion of the reserve tied to the work the shop actually performed. A qualifying HPT, combustor, or HPC restoration should draw against the reserve assigned to those modules, not automatically against the entire engine reserve account. The objective is not to deny legitimate reimbursement. It is to prevent reserve leakage and preserve funding for the modules still carrying future maintenance exposure.

A module-level reserve allocation also functions as a risk-control framework. It helps keep collected reserves aligned with the engine's actual maintenance condition, module-level life consumption, and likely future cash events.

What This Means for Operators, Lessors, Investors, and MRO Providers

For operators: Plan removals around the next binding module, not simply whole-engine time since shop visit. A targeted module repair may preserve availability, but it may also leave the next limiting event unresolved. The right question is not simply how quickly the engine can return to service. It is whether the selected workscope buys enough reliable utility for the aircraft’s planned mission and remaining fleet life.

For lessors: Assess redelivery exposure by module condition, LLP life, records, and next-event risk. A limited module intervention may be commercially sensible for an operator, but it may not create the same residual-value support as a deeper performance restoration or LLP replacement.

For investors: Engine-type exposure is not the same as maintenance-condition exposure. Two engines of the same type, thrust rating, and age may have materially different values if their module histories and LLP positions differ. Underwriting by engine family alone averages assets that may not be economically comparable.

For MRO providers: Demand planning should focus on scope mix, not just shop-visit count. Minimum repairs, performance restorations, overhauls, LLP replacement, parts availability, and turnaround requirements create very different capacity burdens.

Bottom Line

The market buys, leases, finances, and often appraises an aircraft engine as one asset. But it does not consume maintenance capital that way.

It deteriorates by module. It enters the shop by module. It recovers utility by module. And it creates future cash exposure by module.

Whole-engine metrics still matter, but they are only the starting point. Time since shop visit, total flight hours, and headline reserve rates do not answer the real economic question.

The better question is which modules have consumed their utility, which modules the shop has restored, which LLP lives remain, and which section is likely to drive the next cash event.

The engine may be traded as a single asset. The economics live in the modules.

Selected Sources and Further Reading

1. Ackert, Shannon. *Engine Maintenance Concepts for Financiers.*