White Paper

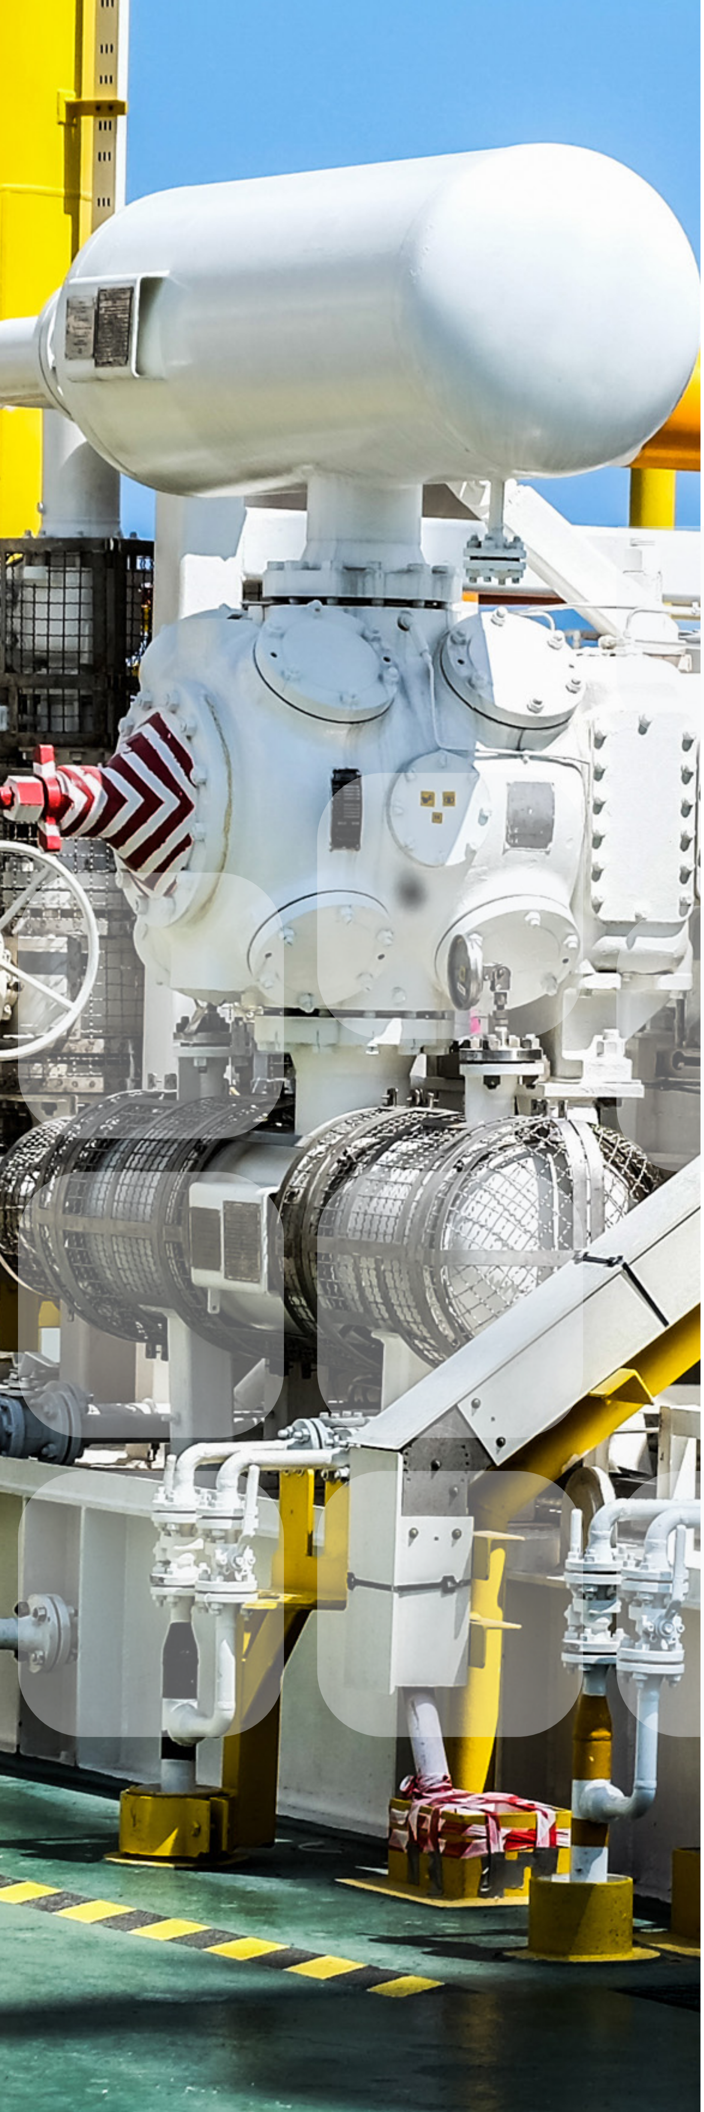

Compressor Reliability Intelligence

A Technical Guide

Reciprocating & Centrifugal Compressors

novity®

www.novity.us

The Compressor Monitoring Gap

Compressors are some of the most critical (and most expensive) rotating equipment in energy operations. When one goes down unexpectedly, the consequences are immediate: lost production, emergency maintenance mobilization, and cascading effects on downstream processes.

Traditional Approaches all Have Shortcomings

Manual vibration routes collect data on a schedule (monthly, quarterly) rather than continuously. Faults that develop between route intervals go undetected until the next visit, or until they cause a trip.

Infrared scanning for valve cap temperatures is standard practice on reciprocating compressors. It works, but it catches valve leaks late. Temperature rise is often a trailing indicator; by the time the IR scan flags it, the valve has already been leaking for days.

OEM alarm systems are configured around fixed thresholds. They trigger when a parameter exceeds a limit, which means they detect that something is already wrong. They don't forecast that something is developing.

Single-parameter monitoring (watching one process signal in isolation) misses the interactions between signals that indicate a developing fault. Incipient issues often show up as subtle shifts across multiple parameters, not a spike in one.

The core problem is threefold. First, lead time: the earlier a developing fault is detected, the more options the maintenance team has to plan the work, order parts, and schedule around production windows. Second, progression visibility: knowing that something is wrong is less useful than knowing how fast it's getting worse and how much operating life remains. Third, actionable precision: a generic alarm tells you to investigate; a specific diagnosis with supporting signal evidence tells you where to look and what to prioritize. Without all three, monitoring creates work without improving decisions.

How Novity Monitors Compressors

Novity's TruPrognostics AI uses a hybrid approach that combines physics-based models with AI. Every model is calibrated to a specific machine type and configured for the data available at that site. Generic anomaly detection flags that something changed; Novity identifies what's failing, why, and how long you have to act.

The Models That Drive the Outcome

The physics models inside TruPrognostics are not general-purpose anomaly detectors. Each one encodes a specific set of thermodynamic or mechanical relationships for the machine type it covers, and each one is designed to surface a specific category of fault.

As an example, on reciprocating compressors, a thermodynamic model takes temperature and compression ratios as inputs and estimates the throw-specific polytropic efficiency for the given compressor configuration (load step, clearance pocket setting, etc. - see Figure 1). These engineering and configuration parameters can be estimated by TruPrognostics AI during calibration.

Changes in efficiency are the primary indicator of mechanical problems in each throw: valve leakage, piston ring wear, packing wear, and similar degradation modes that show up as efficiency loss before they show up anywhere else. Given an estimation of what the efficiency should be, another model predicts the discharge temperature expected for each throw, using the suction temperature and compression ratio as inputs. A measured discharge temperature that is greater than the predicted discharge temperature is a key indication for valve issues. A separate stage capacity model uses compression ratio, speed, and the machine's volumetric geometry to estimate mass flow through each stage. When actual capacity deviates from the model's expectation, it points to internal leakage, interstage dropout, or leaking dump and recirculation valves. These are the types of problems that reduce throughput but may not trigger a temperature alarm for days or weeks. This is just one example.

FROM PHYSICS MODEL TO TIME-TO-ACTION

Shown: discharge temperature model on a reciprocating compressor throw

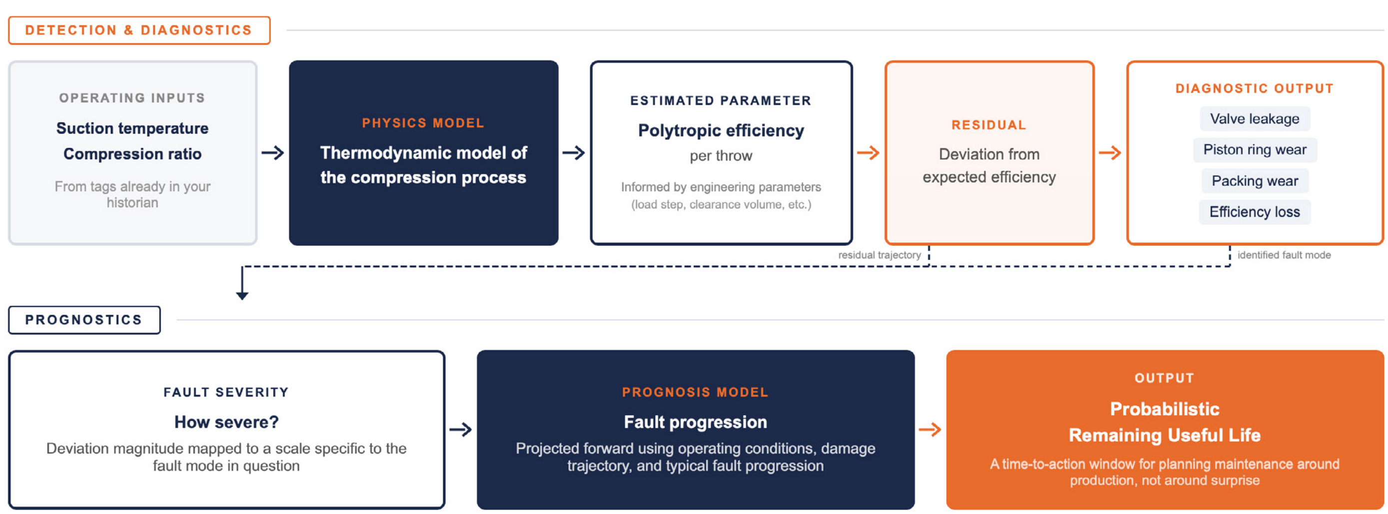

One model of many. TruPrognostics runs multiple models concurrently on each monitored machine. Each model estimates a different performance parameter from the available operating data. When the estimated parameter deviates from the measured value, the residual feeds into the diagnostic AI. The pattern of residuals across models identifies the fault mode; the magnitude and rate of change inform the prognosis. The discharge temperature model shown here is one example. Separate models for stage capacity, interstage pressures, stage gas horsepower, and other parameters provide additional residuals that support differential diagnosis and broader fault coverage.

Figure 1 – Explaining Novity TruPrognostics

TruPrognostics includes a broad library of physics models for each machine type — covering interstage pressures, stage gas horsepower, and many others — that work together to support differential diagnostics across a wide range of fault modes.

Centrifugal compressor models follow the same architecture with different physics. Stage-level efficiency, head-flow characteristics (i.e. compressor “speed curves”), and surge margin replace throw-level volumetric calculations, but the diagnostic logic is the same: the model predicts what the machine should be doing given its current operating point, the residuals reveal deviations from expected behavior, and the pattern across residuals identifies the fault. Process-domain faults like efficiency loss, compression loss, leakage, flow obstruction, and fouling are all detectable at the Base tier (see the Novity Data Tier Model for more details) because the models work with the pressure, temperature, flow, and speed data most operators already collect.

STEP 1

Physics Model Predicts Expected State

A machine-specific physics model predicts the expected value of every monitored tag under current operating conditions and compressor configuration (e.g. load step, clearance volume, etc.).

→

STEP 2

Measured State Deviates from Prediction

The difference between measured and predicted values becomes a residual for each signal. Healthy operation keeps residuals near zero.

→

STEP 3

Pattern and Magnitude Identify the Fault

The pattern of residuals is classified using model-based diagnostics: the measured pattern is compared against known fault patterns and a list of the most likely diagnoses that match is produced.

HOW A RESIDUAL IS FORMED

Single channel shown: discharge pressure

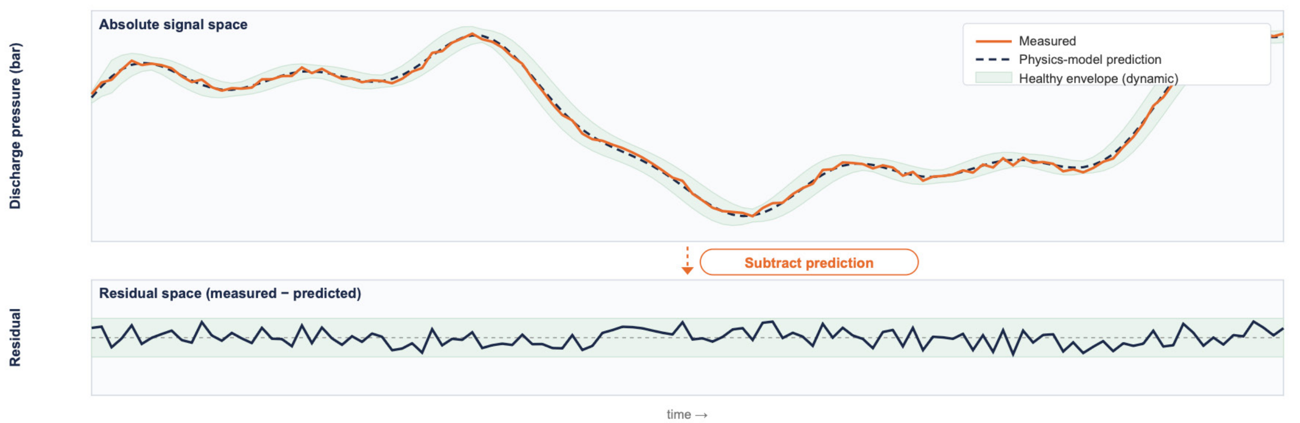

What this shows. The physics model predicts the expected value of each tag from the current operating condition, so the "expected" line moves with the machine. The healthy envelope around it expands and contracts with operating point — wider during transients, tighter at steady load. The residual below is what remains after the model's expectation has been subtracted. Because the dynamic part is already removed, the healthy band in the residual grid below looks flat — but it is not a threshold on the raw signal.

Figure 2 - Residuals Overview

From Models to Decisions

Figure 1 traces one model from operating inputs to time-to-action window. Figure 2 shows what this looks like in practice on a single channel: the physics model’s prediction tracks the machine’s operating point, and the residual is what remains after that prediction is subtracted. Because load, speed, and ambient conditions are already accounted for, the residuals are insensitive to normal operating variation. This is the basis of **detection**: the platform identifies meaningful departures from expected behavior without generating the nuisance alarms that fixed high/low limits produce every time a machine ramps, recycles, or changes duty.

Detection alone does not tell you what to do about the

Reading the figure. Each row is one machine state; each column is a process-domain residual channel. The shaded band marks healthy variation. A valve leak drives coordinated drift across pressure and discharge temperature residuals. Intercooler fouling drives a distinct thermal signature — interstage and discharge temperatures drift positive while pressures stay in band. TruPrognostics recognizes each fault by its full multivariate signature, not by thresholds on any single channel.

Figure 3- Fault diagnosis signatures

problem. **Diagnostics** requires looking at the pattern of residuals across multiple models and tags at once (see Figure 3). Each fault mode has a characteristic signature grounded in the physics of the machine. A valve leak shows up as a coordinated drift across pressure and discharge temperature residuals. Intercooler fouling shows up as a temperature-only signature — interstage and discharge temperatures climb while pressures stay normal. The platform matches the observed pattern against its library of fault models and reports the specific fault mode by name, rather than a generic "anomaly."

Once the fault is identified, **prognostics** tracks how fast the deviation is growing, normalizes it against the asset's healthy and failed states, and projects a probabilistic time-to-action window — so maintenance can be planned around turnarounds, spare availability, and production schedules rather than driven by surprise.

The fault coverage tables later in the white paper show the full list of fault modes the platform covers at each data tier. The models described here are what make that coverage possible — each row in those tables traces back to one or more physics models generating the residuals that drive **detection**, **diagnostics**, and **prognostics**.

Two Analytical Domains, One Platform

Compressor health depends on two fundamentally different types of analysis. Process-domain monitoring — watching pressures, temperatures, flows — catches faults that manifest as changes in thermodynamic or hydraulic performance: efficiency loss, capacity degradation, fouling, leaks. Frequency-domain monitoring — analyzing vibration spectra, HF cylinder pressure waveforms — catches mechanical and bearing faults that show up as changes in spectral signatures long before they affect overall vibration levels or process parameters. Each domain sees things the other misses. Most operators use both for good reason, but typically through separate systems, separate vendor relationships, separate dashboards, and separate alert streams that don't talk to each other.

Process-Domain Analytics

Uses PI data like pressures, temperatures, flows, to detect, diagnose and prognose faults that are observable through this type of data, like leaks, fouling, efficiency losses, and similar.

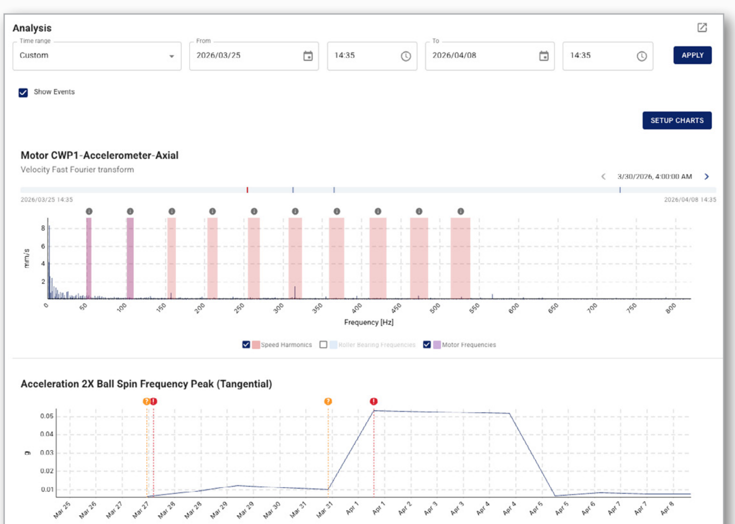

Frequency-Domain Analytics

TruPrognostics AI can process and visualize raw waveform data, such as vibration data for mechanical and bearing fault diagnosis. Rather than using separate vibration and process-domain monitoring, Novity unifies it under one single service.

TruPrognostics's frequency-domain analytics support vibration, current, and pressure waveforms, each of which may be appropriate for different machine types and circumstances.

Novity's TruPrognostics AI covers both in a single platform. At the Base tier, you get full process-domain coverage. When you add high-frequency sensors (vibration, HF cylinder pressure, motor current), the same platform expands into frequency-domain analysis. The models are designed to fuse insights across signal types, so a developing fault that shows subtle process shifts and early vibration signatures gets caught earlier than either domain would catch it alone. For operators, this means one platform, one set of workflows, and one integrated view of compressor health across both analytical domains.

Process-Domain Analytics

Uses PI data like pressures, temperatures, and flows to detect, diagnose, and prognose faults that are observable through this type of data: efficiency loss, compression loss, fouling, leaks, and similar. This is the process engineer's intelligence, and the physics models described earlier in this paper operate primarily in this domain.

At Base tier, process-domain models cover a wide range of fault modes from efficiency loss and compression changes to fouling, leaks, and strainer clogging using only data most operators already collect. Because these models operate on PI/historian data rather than high-frequency sensors, they provide fleet-wide coverage from day one without additional instrumentation.

Frequency-Domain Analytics

TruPrognostics uses **FFT-based analysis** of vibration waveforms to diagnose and prognose mechanical and bearing faults. The benefit over process-domain monitoring for these fault types is twofold: spectral signatures of mechanical degradation typically appear well before overall vibration levels rise enough to trigger a conventional alarm, and the spectral pattern identifies the specific fault mode (bearing defect frequency, imbalance, misalignment, looseness) with higher diagnostic precision and confidence than process data alone can provide. This is the vibration analyst's intelligence, encoded in models and algorithms rather than dependent on a specialist reading a spectrum plot.

The platform also supports **high-frequency in-cylinder pressure analysis** for reciprocating compressors. HF pressure waveforms reveal valve leaks and other cylinder-level faults significantly earlier than discharge temperature or efficiency changes make them visible through process data. This gives operators a second, independent path to valve diagnostics.

In addition to frequency-based analysis, TruPrognostics performs pressure-volume analysis without the need for crank-angle measurements. See this article published in the Oil & Gas Journal for more details.

For motor-driven compressors, TruPrognostics includes **electrical signature analysis (ESA):** spectral analysis of motor current waveforms. ESA can detect rotor bar defects, eccentricity faults, and mechanical load anomalies. Because motor current reflects both electrical and mechanical conditions, ESA provides an independent diagnostic path for faults that would otherwise require vibration sensors to detect.

The workflows and platform interface are the same across both domains. What changes is the underlying class of models, the algorithms, and the visualizations. An operator or reliability engineer works from one consolidated view regardless of whether the insight came from a process residual, a vibration spectrum, or an HF pressure waveform.

The Novity Data Tier Model

Not Every Machine has the Same Instrumentation

Novity's models are designed to work at 3 tiers, each defined by what data is available:

Tier	Data Required	What it Enables
Base	PI / time-series process data (pressures, temperatures, flow, speed) already collected by your SCADA/historian system.	Process-domain diagnostics and prognostics. Detects efficiency loss, compression changes, and some bearing conditions. Works with data you already have.
Plus	Base data + high-frequency vibration or pressure sensors (raw waveform data) on select machines.	Precise mechanical and bearing diagnostics on critical machines. Valve leakage diagnosis on reciprocating compressors. Oil whip/whirl detection on journal bearings.
Premium	Multiple high-frequency sensors and modalities (e.g., vibration + high-frequency cylinder pressure + crosshead/piston rod sensors).	Maximum diagnostic precision: individual valve identification, crosshead wear, rider band wear, piston ring leakage, and more precise efficiency prognostics.

Day-One Value from Existing Data and a Built In Growth Path

Most operators already collect the data Novity needs to start. Suction pressures, discharge pressures, temperatures, flow rates, speed are likely sitting in your historian right now. The Base tier puts that data to work immediately, delivering process-domain diagnostics and remaining useful life estimates for the fault modes your process data can resolve.

That includes efficiency degradation, compression loss, bearing wear trends, and intercooler fouling, flow obstruction, and strainer clogging. No new sensors, new infrastructure, or 6-month implementation project. Connect your cloud-hosted PI data, and Novity's models start generating actionable intelligence on the compressors you're already running.

The upgrade path is purely additive. Every Novity model ships pre-built to support all 3 tiers. When you add a vibration sensor to a compressor already monitored at Base, you don't buy new software or wait for a new integration. The TruPrognostics platform automatically has the requisite models pre-built.

That means you can start lean and scale strategically:

- Start at Base across your full compressor fleet with the data you already have. Get immediate visibility into efficiency loss, compression changes, and bearing health.
- Move to Plus on your highest-criticality machines by adding high-frequency vibration sensors to each unit. This is where certain mechanical fault diagnoses and bearing fault diagnoses become available. It's the biggest capability jump per sensor dollar spent.
- Scale to Premium on the machines where maximum precision justifies the instrumentation investment.

Each step up the tier ladder delivers incrementally more precise diagnostics and longer prognostic lead times. And because the platform is cloud-native with automatic updates, every machine (at every tier) always runs the latest models.

Fault Coverage Reference

Reciprocating compressors have several failure-prone subsystems: valves, pistons, rings, rod packing, crossheads, and journal bearings. The table below shows what Novity can detect, diagnose, or predict for each fault mode at each data tier.

	Premium	Plus	Base
Compressor			
Loss of efficiency	P	P	P
Individual valve leakage	P	A	A
Suction valve Leak	P	P	D
Discharge valve leak	P	P	D
Valve excessive power losses	D	D	D
Piston ring/rod packing leak	D	D	
Rider band wear	P		
Loose piston	D	A	
Loss of compression	D	D	
Stage leak/compressor leak	D	D	D
Loss of efficiency	P	P	P
Individual valve leakage	P	A	A
Frame and Running Gear			
Main bearing wear	P	P	A
Oil whirl/oil whip/unbalance/misalignment	D	D	
Crosshead bushing wear	D	A	A
Loose crosshead	P	A	A
Frame looseness	P	P	A
Frame loss of lubrication	D	D	D
Lube oil filter clogging	P	P	P
Main bearing wear	P	P	A
Auxiliary Systems			
Suction piping restriction	D	D	D
Scrubber fouling	D	D	
Excessive dropout/liquid in gas stream	D	D	A
Intercooler/aftercooler fouling	D	D	D
Strainer clogging	P	P	P
Cooler fan failure	P	P	A
Cooler fan bearing wear	P	D	A
Suction control valve fouling	D	D	D

P = Prognostics
D = Diagnostics
A = Detection

	Premium	Plus	Base
Engine / Driver			
Engine overspeed/underspeed	D	D	D
Engine loss of lubrication	D	D	D
High engine vibration	P	P	A
Engine unbalance/misalignment	D	D	A
Engine oil filter clogging	P	P	P
Engine overheating	D	D	D

P = Prognostics
D = Diagnostics
A = Detection

What This Means in Practice

At Base tier, using only your existing PI data, Novity can already predict throw-level efficiency loss, a wide range of compressor fault modes, and bearing wear with remaining useful life estimates. It can also detect (flag as abnormal) imbalance and misalignment conditions. This covers the most common efficiency-related compressor issues without any new sensors.

At Plus tier, adding a single high frequency vibration sensor enables high-precision diagnosis for a wide range of bearing faults modes, as well as several compressor faults modes. This is the biggest capability jump per sensor added.

At Premium tier, with HF cylinder pressure and additional vibration/position sensors, Novity can identify individual valve leaks, diagnose crosshead and rider band wear, and detect loose pistons. This level of granularity is typically reserved for the most critical compressors where maximum diagnostic precision justifies the instrumentation investment.

Supported Data Inputs

Base – PI / Time-Series
Suction & discharge pressures
Interstage pressures
Discharge flow rate
Rotational speed
Cylinder discharge temperatures
Suction temperature
Intercooler / aftercooler temps
Lube oil pressure & temperature
Frame vibration (Vrms)
Packing vent pressure & temp
+ 15 additional optional tags

Plus – Base + Minimal HF Sensors
All Base tags, plus:
High-frequency frame/engine vibration *OR*
High-frequency cylinder pressure

Premium – Multiple HF Sensors
All Plus tags, plus:
Crosshead vibration (HF, 3-axis)
Piston rod drop (HF, 3-axis)
Individual valve cap temperature

Centrifugal Compressors: Fault Coverage by Tier

Fault Mode	Premium	Plus	Base
Compressor			
Loss of Efficiency (by stage)	P	P	P
Loss of Compression (by stage)	P	P	P
Leakage (by stage)	D	D	D
Flow Obstruction	D	D	D
Frame Looseness	D	D	D
Intercooler Fouling	D	D	D
Suction Strainer Clogging	D	D	D
Journal Bearing			
Bearing Wear	P	P	P
Oil Whip	D	D	
Oil Whirl	D	D	
Imbalance	D	D	A
Misalignment	D	D	A
General Looseness	D	D	
AC Motor (if motor-driven)			
Winding Faults	D	D	
Broken Bar	D	D	A
Eccentricity Fault	D	D	A
Overload	D	A	
Mechanical Wear	P	P	A
Motor Control Diagnostics	D	D	D
Winding Insulation Degradation	P		

P = Prognostics
D = Diagnostics
A = Detection

What This Means in Practice

Centrifugal compressors have strong Base-tier coverage. Because the dominant centrifugal failure modes (efficiency loss, compression loss, leakage, flow obstruction, fouling) are process-domain faults, they're detectable using the pressure, temperature, flow, and speed data most operators already collect. You can monitor the health of your centrifugal fleet on day one without adding sensors.

Where Plus and Premium tiers add value is in mechanical diagnostics: bearing instability, looseness, and (for motor-driven units) electrical fault diagnosis. If a centrifugal compressor is critical enough to justify a vibration sensor and motor current measurement, Plus or Premium tiers adds significant mechanical fault coverage.

Case Evidence: Compressor Predictions in the Field

The fault coverage tables show what Novity's models can do in theory. These cases show what they've done in practice.

Cylinder Fault: 4 to 6 Weeks of Lead Time

Customer: U.S. oil and gas supermajor

Machine: 7,800 HP reciprocating compressor

Novity identified a failure trend in June 2024, projecting a fault window of July 10 to July 27. The failure occurred on July 17. That's 4 to 6 weeks of advance notice on a machine where unplanned downtime is measured in hundreds of thousands of dollars per day.

After the compressor was repaired and returned to service, the asset health score reverted to 100% with no fault prognosis present. The model correctly reflected the repair.

1

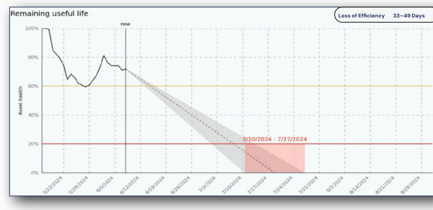

2

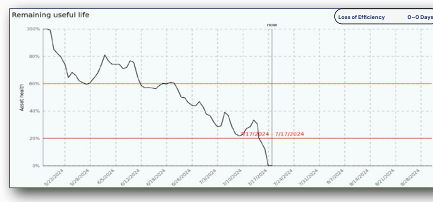

3

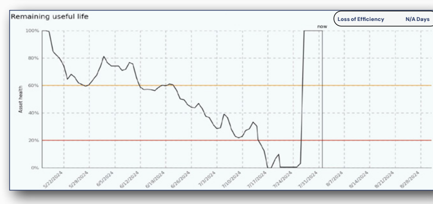

Catastrophic Failure Avoided

Customer: U.S. oil and gas operator

Machine: 4,000 HP reciprocating compressor (multiple units)

Novity flagged a loss of performance in a 3rd-stage cylinder. Temperature readings showed nothing abnormal, so the operator kept the compressor running while monitoring it closely. 3 days later, elevated frame vibration appeared. When the team opened the machine, they found the 3rd-stage piston had become loose. With less than a day of remaining margin, a catastrophic failure and several months of downtime were avoided.

This is a case where traditional monitoring (temperature) missed the issue entirely. The combined signal analysis from Novity caught what a single-parameter approach couldn't.

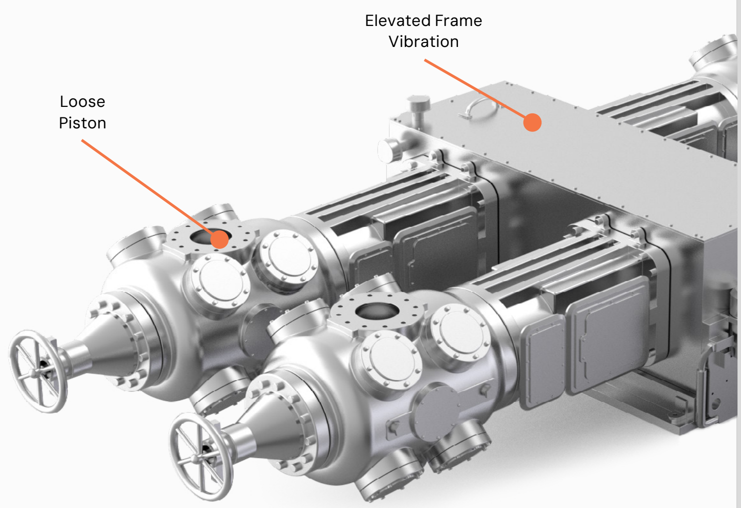

40+ Days Earlier Than OEM Alarm

Customer: U.S. oil and gas supermajor

Machine: Reciprocating compressor

Novity detected a loss of efficiency more than 40 days before the OEM's built-in alarm system would have flagged it. The model provided diagnostics with multiple plausible fault modes (suction valve leak, discharge valve leak, excessive valve power losses), each with supporting signal evidence.

40+ days of additional lead time changes the conversation from “emergency shutdown” to “planned maintenance during the next scheduled window.”

Getting Started

Technology is only half the question. Reliability and operations leaders also need to know what deployment looks like, how fast they can see results, and what their team needs to do.

Start with a Focused Evaluation

Novity can run an offline assessment on a small set of compressors using historical data you already collect. You export and share historical tag data for a few machines — no live connections, no hooks into your system, no integration work needed. Novity configures the TruPrognostics models, runs them against your data, and delivers a report covering detected conditions, diagnostic findings, and prognostic estimates where the data supports them. No operational risk, no commitment beyond sharing the data. Typical turnaround is 2 to 4 weeks.

Your team sees real findings on your actual equipment before any deployment decision. What faults are developing, how fast, and what the data supports for each machine.

What you Need to Get Started

All available tags associated with each compressor you want to evaluate. Suction and discharge pressures, temperatures, flow, and speed are particularly important, but the more complete the tag list the better the evaluation — additional tags often improve model accuracy and enable broader fault coverage. If you want to include high-frequency data (vibration, cylinder pressure) for specific machines, Novity can incorporate that into the evaluation as well.

How Full Deployments Work

Novity connects to your cloud-hosted data (PI, historian, or SCADA data made available via API or standard integration). TruPrognostics models are configured for your specific compressor types and operating conditions. Adding a new machine is a matter of providing tag access — Novity configures the models to the available data and onboards new machines in a matter of days. No new model development required. Machines with varying instrumentation levels coexist on the same platform, each getting the best diagnostics and prognostics their data supports.

Diagnostics, prognostics, and maintenance recommendations are delivered through the TruPrognostics platform, accessible via web browser. Updates are automatic. As models improve, your site runs the latest version without any action on your part. Novity routinely achieves fault detection true positive rates in excess of 90% across production deployments, with low false positive rates. Fault mode coverage is documented per model version and data tier, so your team knows exactly what to expect from each level of instrumentation.

Starting from existing tags and growing capabilities from there is the fastest and most cost-effective path forward.

Talk to Novity About Your Compressor Fleet

Whether you're running reciprocating compressors, centrifugal compressors, or both, we can walk through what Novity would look like on your equipment, with the data you have today.

Contact Us

info@novity.us

Website: novity.us/contact