

Supporting Real-Time Quality Insights Through HL7® FHIR®

A Federal Policy and Industry Roadmap to
Modernize Digital Quality Measurement
Reporting for ACOs

Ryan Howells

Mark Marciante

Sarah Hudson Scholle

Daniela Simpson

OCTOBER 2026



Contents

Executive Summary.....	3
The Next Step in Digital Quality Reporting	6
Current CMS Regulations and the policy opportunity.....	7
Why Is Change Necessary?	7
ACO Experience with Certified EHR Technology and FHIR	8
A Target-State Architecture.....	11
The Current Workflow for ACO Digital Quality Reporting.....	11
The Workflow That Should Be Possible	11
A Three-Part Framework for FHIR Infrastructure.....	12
1. FHIR Conformance	12
2. Data Quality.....	14
PIQI: Moving Beyond “The Data Are There”	14
3. API Performance.....	15
Proposed API Performance Requirements	15
The Cost of Full Population Digital Quality Measurement.....	17
A New Policy Roadmap Built On FHIR.....	19
Other Obstacles To Digital Quality Measurement For ACOs.....	21
Conclusion	22
About the Authors	23
About Leavitt Partners	23
About Wakely	23

EXECUTIVE SUMMARY

Quality measurement is the mechanism that the Centers for Medicare & Medicaid Services (CMS) uses to assess whether Accountable Care Organizations (ACOs) are improving care for the Medicare beneficiaries they serve. Unlike individual providers, ACOs must make that assessment across a population that receives services from multiple clinicians, practices, hospitals, and technology platforms. Health care quality reporting, however, is labor-intensive, complex, and costly for ACOs. Much of that burden comes from how quality data are acquired, exchanged, validated, and refreshed. In an effort to move toward digital quality measurement, the agency has focused on a quality reporting pathway that requires the use of electronic health records (EHRs) and measure specifications built on a quality-specific data model.

There is a better approach. Modern, standards-based infrastructure based on currently regulated tools could reduce the burden of quality measurement and ensure that quality insights become part of ACO's population health strategy. The industry now has many of the technical building blocks needed to improve how quality is measured, validated, and reported, including certified EHR technology and standardized application programming interfaces (APIs), requirements for using HL7® Fast Healthcare Interoperability Resources (FHIR®), and the growth of health information networks and investment by clinical organizations in data repositories, analytics platforms, and population health tools. These policies and infrastructure have advanced to create the ability for ACOs to manage the health of their population more effectively.

Consequently, a growing disconnect exists between the technologies ACOs use to improve outcomes and those CMS expects them to use to report those outcomes. The greatest opportunity lies in aligning quality reporting with modern interoperability infrastructure and ensuring that it delivers timely, meaningful data. CMS recognizes this need and is moving towards this approach. In the [CY 2027 Physician Fee Schedule \(PFS\) proposed rule](#), CMS issued a request for information on a plan for moving to digital quality measurement supported by FHIR as well as other changes to the Medicare Shared Savings Program (MSSP) rules that show their desire to move to FHIR-based infrastructure.

Still, further regulatory changes are necessary to make this architecture possible without requiring it. We recommend that CMS give ACOs the flexibility to use any standards-based FHIR API infrastructure that aggregates data from participating providers and other permitted sources and makes the information available for quality measurement and other uses. Such flexibility, however, must be accompanied by clear requirements to ensure the infrastructure can support trusted, auditable quality reporting at scale. Any standards-based FHIR API qualifies for this purpose, including clinical APIs and claims APIs, as well as those operated by health information exchanges (HIEs), certified data aggregators and population health platforms. Document-based exchange, including the Continuity of Care Document, does not.

A practical three-part framework for evaluating FHIR APIs is necessary and should focus on the following considerations:

- **Does the API conform?** The endpoint should demonstrate that it implements the required FHIR resources, profiles, interactions, authorization mechanisms, and other technical requirements through use of the Office of the National Coordinator's (ONC's) Inferno Test Kit.
- **Are the data good enough?** The information returned by the API should contain the data elements and quality characteristics the quality reporting use case requires. The Patient Information Quality Improvement (PIQI) Framework recently balloted at HL7 could provide a foundation for that baseline.
- **Can the API support performance required for population health?** The endpoint should respond quickly and reliably to queries covering a large population, not just a single patient record. If an API facade sits in front of a source system, both should meet the requirement.

As CMS moves to population-level health quality reporting, a new category of recurring infrastructure also needs to be considered—cloud computing. Moving to population-level reporting is the right policy direction because it aligns quality with population health management. FHIR is the only standards-based mechanism mature enough to make full population reporting operationally feasible. This shift, however, carries a cloud infrastructure cost that is unaccounted for in the current program design. Modeled on reserved capacity across storage, compute, and egress, the cost runs \$1.37 to \$1.73 per covered life per year, scales with population size, and recurs annually regardless of program outcome. It falls disproportionately on organizations too small to negotiate the volume pricing that cloud vendors and calculator vendors reserve for their largest customers.

CMS can make this modern architecture possible (though not yet required) through several changes, including:

- **Redefine certified EHR technology (CEHRT) at 42 CFR § 414.1305** of the Merit-based Incentive Payment System (MIPS) provisions in the Medicare Access and CHIP Reauthorization Act of 2015 (MACRA) to include any API certified under the ONC Health IT Certification Program to the standards adopted at 45 CFR 170.215.

- **Create a new Medicare digital quality measures (Medicare DQMs) data collection type as a voluntary pathway** for quality reporting that uses Standardized FHIR-based APIs in 42 CFR § 414.1305 and 42 CFR § 425.512. This change will allow ACOs to meet the quality performance standard if they report quality data using Medicare DQMs. Allowing pay for reporting for this option would be an incentive for organizations to attempt this option and would give CMS valuable data and the opportunity to work through implementation challenges, without requiring organizations to put millions of dollars in shared savings at risk.
- **Add an option for meeting MSSP CEHRT use requirements at 42 CFR § 425.507 by reporting objective results for conformance, data quality, and performance of the FHIR APIs** used for reporting of at least one of the five ACO-reported measures in the Alternative Performance Pathway (APP) Plus measure set. ACOs would report “yes” for each measure by showing that the APIs passed conformance testing for the required FHIR standards through ONC’s Inferno Test Kit, met a data quality baseline such as the HL7 Patient Information Quality Improvement (PIQI) Framework, and met performance requirements for population-level queries. This option would tie the CEHRT requirement to the infrastructure ACOs use for quality reporting and would only be available to ACOs electing the Medicare DQM pathway.
- **Implement steps to relieve the cost of cloud computing** necessary for Medicare DQMs by facilitating shared infrastructure or finding ways to extend favorable volume pricing to entities undertaking Medicare DQM collection.

Ongoing work by CMS, ONC, HL7, the National Committee for Quality Assurance (NCQA), and the private sector has demonstrated the value of a modern, standards-based approach to quality reporting. The recommended changes would let ACOs report quality performance from the same FHIR infrastructure they use to manage care. Quality insights would then be available for care delivery and improvement, and reporting would stay accountable and auditable.

THE NEXT STEP IN DIGITAL QUALITY REPORTING

Much of the burden of quality measurement comes from how quality data are acquired, exchanged, validated, and refreshed. A published analysis¹ of hospital quality reporting supports this perspective. When reporting is decomposed by activity, cost concentrates in collecting and validating data, while measure calculation is a comparatively minor expense. Burden also grows with the number of reporting programs that require overlapping metrics that are collected and submitted in different ways.

The first phase of efforts to achieve health care interoperability focused on making information available. The next phase must focus on making that information useful. Many of the technical building blocks necessary to improve how quality is measured, validated, and reported are now available. ONC established requirements for certified Electronic Health Record technology (CEHRT) and standardized application programming interfaces (APIs) for patient and population services, and Inferno provides executable testing for HL7® Fast Healthcare Interoperability Resources (FHIR®) implementations. The Centers for Medicare & Medicaid (CMS) incorporated FHIR into its own interoperability rules under the 21st Century Cures Act. At the same time, health information networks expanded nationwide, and health care organizations invested heavily in clinical data repositories, analytics platforms, and population health tools. These investments have made it increasingly possible to manage, validate, and reuse clinical data across systems.

With the infrastructure increasingly in place, the policy question now is whether CMS will allow that infrastructure to support the organizations responsible for managing care across populations. This issue is especially relevant to accountable care organizations (ACOs) that participate in the Medicare Shared Savings Program (MSSP). CMS is asking ACOs to accept greater responsibility for population health, quality improvement, and financial accountability.

The ACO has two related responsibilities: Manage the care of the population and demonstrate value.

The technology supporting those responsibilities should not be artificially separated. If an ACO has invested in a population-level clinical data repository, analytics environment, and FHIR infrastructure, it should be able to use that infrastructure to support quality reporting to CMS.

¹ Saraswathula A, Merck SJ, Bai G, et al. The Volume and Cost of Quality Metric Reporting. JAMA. 2023;329(21):1840-1847. doi:10.1001/jama.2023.7271. Time-driven activity-based costing study of 162 externally reported inpatient quality metrics at a single academic medical center, at an annual cost exceeding \$5 million in personnel time plus more than \$600,000 in vendor fees.

CURRENT CMS REGULATIONS AND THE POLICY OPPORTUNITY

CMS regulations specify the quality measures that ACOs must report and establish the framework for determining whether an ACO satisfies its quality performance standard. Through MSSP, the agency expects ACOs to manage population health across multiple EHRs, improve quality and assume accountability for outcomes. The technical submission requirements, however, remain tied to EHR-based quality reporting mechanisms rather than to broader recognition of modern, standards-based population health infrastructure.

CMS defines the methods that ACOs use to submit quality measures for MSSP and other Advanced Alternative Payment Methodologies (APMs) (42 CFR § 414.1305). CMS has positioned electronic clinical quality measures (eCQMs) as the pathway to digital quality measurement. The flaw in the eCQM approach is the reliance on quality-specific data architecture and submission from EHRs. Clinical concepts are represented through a quality-specific data model, the Quality Data Model (QDM), packaged for submission through a quality-specific document format, the Quality Reporting Data Architecture (QRDA), and dependent on CEHRT.

CMS defines quality measures for ACOs under 42 CFR §425.508 of MACRA and has explicit authority to audit and validate submitted data, which creates a practical policy opportunity. CMS can maintain its quality measures, data completeness requirements, and audit and validation authority. It simply needs to allow a different technology architecture to satisfy those requirements. That modern architecture should build on FHIR-based interoperable data to create a qualified, standards-based environment that ACOs can use to aggregate data from participating providers and other permitted sources for quality reporting, as well as other use cases.

Why Is Change Necessary?

Change is needed for two reasons: 1) the presumed role of the EHR, and 2) the reliance on a quality-specific data model. The EHR is one of the most important technologies in American health care delivery, yet it was never designed to serve as a population health platform. It supports clinical care and serves as the system of record for medical encounters, but it is just one technology used to manage a patient population. That distinction matters because modern ACOs operate across a much broader information environment than any single clinical application can support.

A typical ACO may need to bring together information from multiple EHR platforms, physician practice systems, hospital systems, HIEs, claims databases, clinical registries, care management platforms, population health systems, enterprise data warehouses, analytics platforms, laboratory and pharmacy systems, and other clinical and administrative data sources.

As the Leavitt Partners [Kill the Clipboard!](#) white paper noted, provider data exchange does not occur exclusively through EHRs. Though EHRs remain critical sources of clinical information, ACOs also depend on vendor applications, clinically integrated networks, value-based care platforms, analytic warehouses, and other systems that aggregate and use EHR data.

That observation has become more important as value-based care has expanded. An ACO is an organization responsible for understanding and improving the health of a defined population—not simply a collection of EHRs. Fulfilling that obligation requires information architecture capable of bringing data together, validating the information, and turning the findings into actionable insights.

The second issue is the use of a quality-specific data architecture. ONC maintains the [United States Core Data for Interoperability \(USCDI\)](#), which defines a core set of health data that all electronic health systems should be able to share. Today, certified FHIR APIs are required to support USCDIv3. Prior to the adoption of USCDI, CMS developed the QDM, which has defined the [USCDI+Quality](#)—a separate data model for quality that includes many data elements excluded from the required core. The disconnect between these models has created unnecessary barriers to the use of FHIR for quality reporting. A [recent Leavitt Partners analysis](#) showed that 49 program-specific measures reported by health systems and facilities across 13 programs can be reported today using current USCDIv3 data. Adding just 10 targeted data elements would enable reporting of another 26 program-specific measures, reaching almost two-thirds of the priority measures.

ACO EXPERIENCE WITH CERTIFIED EHR TECHNOLOGY AND FHIR

Several ACOs testing FHIR®-based quality reporting for MSSP—as well as one that is already reporting using FHIR® APIs—identify benefits but note a consistent set of limitations.

eCQM costs grow with each EHR instance and every reporting program. FHIR costs do not.

Multiple surveys of leading ACOs have confirmed that QRDA-based output serves only the purpose of submission. Each additional instance and each additional program adds cost. Vendors price the calculator per provider and deliver a single annual calculation.

For example, **ACO A**, a large organization that must connect to 171 unique EHR instances, pays vendors to establish and maintain each data connection to obtain quality data. Charges range from \$500 per connection with a \$600 annual maintenance fee across hundreds of instances, to one-time charges of \$3,300 and \$5,000, to \$15,000 for a QRDA file covering a two-provider practice. FHIR connectivity, in contrast, grows with the number of connections and regulated FHIR APIs carry no additional charge.

FHIR-based reporting is more cost-effective. **ACO B**, a large health system with two MSSP ACOs, reports across multiple Epic instances and approximately 20 external EHRs. To calculate eCQMs, it stood up a separate instance of its EHR whose only function is to aggregate data from the others for quality reporting but has not yet produced a complete measure calculation. Patient matching, combining data across instances, and incorporating external data sources remain unresolved. ACO B estimated reporting cost of each approach for a single run of four MSSP measures covering breast cancer screening, colorectal cancer screening, depression screening, and diabetes control.

Table 1. Estimated Cost of Reporting Four MSSP Quality Measures Using eCQMs vs. FHIR for One ACO

	eCQMs	FHIR-Based DQMs
Technical Resources	\$665,000	\$100,000
Business Resources	\$285,000	\$55,000
Measure Calculator	No incremental costs ²	\$250,000
Total Estimate (Annually)	\$950,000	\$405,000
eCQMs: Electronic clinical quality measures		
FHIR-Based DQMs: Proposed digital quality measure reporting based on regulated FHIR endpoints		

As **Table 1** demonstrates, ACO B estimated that eCQM reporting costs would be more than twice its projections for FHIR-based DQM reporting (\$950,000 versus \$405,000, respectively) because the same underlying infrastructure also meets the ACO's quality reporting obligations under value-based contracts with multiple Medicare Advantage plans. The cost for FHIR DQM is incremental rather than new, with annual technical resources expected to cost \$100,000, well below the \$665,000 required for eCQMs.

Calculating Medicare clinical quality measures (CQMs) from HL7 FHIR data is feasible, and the benefits extend beyond the submission. An ACO with 11 separate EHRs and 29 practices, including independent practices, reported that: 1) once established, the connectivity and calculation workflow created a scalable foundation for future reporting; 2) engagement in quality improvement increased across the organization and among clinicians; and 3) the reporting effort itself seems to be driving greater standardization across workflows.

² Measure calculation is included in existing CEHRT licensing. The work required to assemble and validate the data that feeds the calculation is not and appears in the technical resources line.

Still, ACOs that are building capacity to use FHIR-based reporting are encountering multiple challenges. ACOs report that the FHIR standard already includes the data elements required for quality measure reporting. The issues stem from how certified systems approach implementation.

- Vendors implement FHIR to pass certification rather than to exchange data. Vendors publish conformant APIs and then direct customers toward proprietary interfaces on the grounds that the proprietary path is faster and better supported. ACOs working across many EHR platforms encounter this challenge in their work with multiple vendors.
- Certified APIs do not perform at population scale. Current configurations require several calls per patient, which produces tens of thousands of calls for a population of 10,000. Bulk requests cannot be scoped to the population volume or the data elements requested, so ACOs revert to patient-by-patient querying, absorbing the call counts and the duplicative data that follow.
- Extraction competes with patient care. EHRs run extraction on the infrastructure that supports clinical operations. Population scale pulls raise demand at the same time clinical systems require real-time responsiveness, and vendors complain about the production impact.
- Certified systems return local codes rather than required terminology standards. Proprietary codes appear across FHIR resources for the diagnoses, procedures, and results, forcing ACOs to build and maintain crosswalks against large value sets. Data behave differently depending on how they are created. Vitals, immunizations, and laboratory results, which systems code automatically, extract reliably. Encounters, health assessments, and procedures, where a person enters the value or no discrete field exists, extract poorly. A procedure recorded as free text does not match a measure engine expecting a Current Procedural Terminology (CPT®) or Systematized Nomenclature of Medicine (SNOMED) code, and the service goes uncounted although it occurred and would have closed the gap.
- Certified systems exclude codes that are required in the API but include them in quality-specific feeds. In some cases, ACOs have found that LOINC and other codes from external systems are not transferred to the EHR at all, but they can be found in the quality feed. In addition, required structured data exist in text rather than the USCDI-related field, where it could be used for reporting.
- Data quality is poor and certification does not test for it. A vendor can pass certification and still hand the ACO data that cannot be used. Certification confirms the API responds, returns the required resources, and handles authorization. It does not test whether the data behind that interface is complete or coded to the required terminology standards. The ACO learns of this omission too late, after the connection is built and data are pulled.

A TARGET-STATE ARCHITECTURE

The target-state architecture is organized around the ACO's ability to support population management and whole-person care. Clinical data are collected from multiple authorized sources through standards-based FHIR APIs because patients receive care across multiple providers and in various settings, each of which may use different technology platforms. In this model, standardization, data quality, and the timeliness of data are critical. FHIR is uniquely positioned to solve this issue, demonstrating how the target architecture changes the workflow.

The Current Workflow for ACO Digital Quality Reporting

CMS's proposed strategy for moving to digital quality reporting assumes an EHR-centered path:

Provider → EHR → Reporting Process → CMS

That workflow can make sense when quality information is created, maintained, and reported primarily from a single clinical application by a single clinician. It becomes much less attractive when an ACO is accountable for thousands of providers using different technologies across different care settings.

The Workflow That Should Be Possible

Modern architecture building on FHIR-based interoperable data would recognize EHRs as essential sources of clinical data while allowing ACOs to operate the population-level infrastructure needed for quality reporting and other uses as follows:

Provider EHRs and Other Data Sources → Standards-Based FHIR APIs → ACO Cloud-Based FHIR Infrastructure → Data Quality Validation and Measure Calculation → CMS Quality Reporting AS WELL AS for Care Management and Quality Improvement

Providers maintain standards-based FHIR APIs, ACOs use authorized transactions to bring data into a qualified population-level environment, and CMS receives auditable quality information through an approved reporting mechanism. In the target-state architecture, data originate from any authorized source exposing a standards-based FHIR API. Examples include EHR APIs, HIE APIs, laboratory and pharmacy systems, claims APIs operated by payers and CMS, clinical registries, and certified data aggregators. The qualifying condition is the interface, not the category of organization operating it. Document-based exchange falls outside the pathway. Each participant exposes information through standards-based FHIR interfaces, creating a common framework for exchange.

This infrastructure could support multiple authorized functions beyond quality measurement, including care coordination, transitions of care, patient access, payer-provider exchange, prior authorization, population health, clinical decision support, and public health.

A THREE-PART FRAMEWORK FOR FHIR INFRASTRUCTURE

For this modern architecture to be effective, the qualified FHIR endpoints need to avoid the problems that ACOs encounter. A practical three-part framework for ensuring that FHIR APIs are fit for purpose needs to address technical conformance, data quality, and API performance.

- 1. Does the API conform?** The endpoint should demonstrate that it implements the required FHIR resources, profiles, interactions, authorization mechanisms, and other technical requirements through use of the ONC-provided [Inferno Test Kit](#).
- 2. Are the data good enough?** The information returned by the API should contain the data elements and quality characteristics the quality reporting use case requires. The [PIQI Framework](#) recently balloted at HL7 could provide a foundation for that baseline.
- 3. Can the API operate at scale?** The endpoint should respond quickly and reliably to queries covering a large population, not just a single patient record. Where an API facade sits in front of a source system, both should meet the requirement.

The framework works only if all three tests are applied together: conformance shows whether the endpoint follows the standard, data quality shows whether the information can be trusted, and performance shows whether the endpoint can operate at population scale. Details for each aspect follow.

FHIR Conformance + Data Quality + API Performance

1. FHIR Conformance

Question: Does the endpoint speak the required language?

FHIR conformance should establish that the endpoint implements the applicable resources, profiles, interactions, authorization requirements, and implementation guides. Inferno already provides a tested mechanism for evaluating FHIR implementations, including the ONC §170.315(g)(10) standardized API certification criterion. ACO-operated endpoints should build on that same principle by demonstrating conformance to the requirements CMS and Assistant Secretary for Technology Policy (ASTP)/ONC identify for population-level quality reporting.

US Core as required in certified FHIR APIs should be sufficient for quality reporting. If a measure requires a data element outside US Core, that element should be added to USCDI through the established process rather than maintaining a parallel data model for quality. A [recent analysis](#) found that 49 program-specific measures across 13 programs can be reported using USCDIv3 data as it stands today and that 10 additional data elements would enable 26 more. Specifying measures against the data that certified APIs already return removes the disconnect between what an ACO can extract and what a measure requires.

Inferno should move from certification testing to operational conformance.

Inferno is already an important part of ONC's certification ecosystem. The current ONC §170.315(g)(10) test kit helps in evaluating whether certified health IT supports key API capabilities, including SMART App Launch, US Core, bulk data, single patient APIs, and multi-patient APIs. ACO digital quality reporting should build on that capability, but it should not require every ACO reporting environment to become a new ONC-certified health IT product. Instead, CMS should recognize Inferno-tested conformance as one condition for qualifying a FHIR endpoint (see **Appendix 1**).

Under this approach, Inferno would support a chain of technical accountability across the ecosystem. Provider endpoints could be tested for conformance before data are ingested. ACO infrastructure could then be evaluated as part of the qualification process before quality results are submitted to CMS.

This approach would preserve technology neutrality. The relevant question is not whether the endpoint is operated by an EHR vendor, an ACO, a cloud provider, an HIE, or another qualified technology partner. The question is whether the endpoint can demonstrate conformance to the standards CMS requires for quality reporting.

2. Data Quality

Question: Is the information fit for the intended use?

FHIR conformance alone does not answer this question. A resource can conform to a profile while still containing incomplete, stale, inconsistent, or otherwise insufficient information. The [HL7 PIQI Framework](#) provides a structured means of evaluating whether patient-centric data are complete, accurate, timely, and reliable enough for a specific quality measurement use case.

For CMS quality reporting, the relevant question should not be whether the data conform to FHIR, but rather whether the data are sufficient to calculate the measure reliably. PIQI offers a potential framework for making that distinction operational.

PIQI: Moving Beyond “The Data Are There”

Digital quality reporting depends on both interoperability and data quality. It is not enough for information to move from one system to another. The information must also be complete, accurate, timely, and usable for the specific quality measurement purpose.

Historically, interoperability has centered on whether a message, resource, code, or API interaction conforms with a required standard. Those questions remain important, but they do not answer whether the data are robust enough to support a reliable outcome.

The [PIQI Framework](#) provides a practical way to make that distinction. PIQI can be used to evaluate whether patient-level data are available, complete, accurate, consistent, current, and reliable enough for the intended use case.

For digital quality reporting, PIQI could help CMS and ACOs move from a narrow conformance question—whether the data are structured correctly—to a more meaningful accountability question: whether the data can be trusted for measurement.

Conformance shows whether the data follow the standard. PIQI can help show whether the data are usable.

3. API Performance

Question: Can the endpoint perform at population scale?

An ACO endpoint should qualify for population-level reporting not simply because it can pass a conformance test, but because it also can operate reliably under real-world conditions, including large patient panels, bulk data requests, concurrent users, authentication overhead, and sustained query volume. The current ONC certification does not consider performance requirements to be essential, and EHRs vendors are allowed to define their own testing criteria.

API performance should therefore be part of the qualification framework. CMS and ONC should treat population-scale reliability as a key requirement for ACO digital quality reporting, along with conformance and data quality.

Performance requirements should apply to the source system as well as the interface. An API facade can meet response time benchmarks by serving data staged in advance while the underlying source system remains the constraint on how current and complete that data is. Laboratory, imaging, pharmacy, and EHR systems all sit behind interfaces in ACO reporting environments. Testing should establish performance at the source and at any facade, so that the benchmark reflects the data the ACO can obtain rather than the speed of the layer in front of it.

Proposed API Performance Requirements

API performance requirements should reflect how FHIR APIs are used in practice. A single-resource read, a search request, a write transaction, and a population-level bulk data request impose different demands on the underlying infrastructure. Hence, CMS and ONC should establish performance categories rather than rely on a single response time standard.

The proposed benchmarks in **Table 2** should be treated as a starting point for industry testing and discussion, not as immediate regulatory mandates. The first objective should be to validate whether these thresholds are realistic, measurable, and useful in real-world ACO environments.

Performance testing should also account for the operating conditions that materially affect API responsiveness and reliability. These considerations, as outlined in **Table 3**, should be evaluated separately from the benchmark thresholds so that CMS and industry can distinguish between the target performance level and the factors that influence whether an endpoint can meet it.

Table 2. Proposed API Performance Requirements

API Transaction	Proposed Benchmark	Use Case
Single-resource read	Average response time: under 1 second	Retrieving an individual FHIR resource, such as a patient, an encounter, an observation, or a medication record
Search request, no operations, first page	Average response time: under 2 seconds	Returning the first page of search results without additional computational operations
Search request with operations, first page	Average response time: under 3.5 seconds	Returning the first page of search results involving filtering, sorting, chained parameters, or other supported operations
API write transaction	Average response time: under 5 seconds	Submitting, updating, or writing a FHIR resource or transaction where permitted
\$everything operation	Average response time: under 10 seconds	Retrieving a comprehensive patient record where the operation is supported
Bulk data request	Average processing time: under 10 seconds per request	Exporting population-level data for quality measurement, analytics, or reporting

Table 3. Additional Performance Testing Considerations

Consideration	Why It Matters
Payload size and number of resources requested	Larger payloads and broader resource requests can materially affect response time, processing burden, and system stability.
Number of patients	Population-level reporting should be tested against meaningful patient panels rather than small synthetic datasets alone.
Concurrent users and concurrent requests	Testing should reflect realistic demand from multiple users, systems, or reporting processes operating at the same time.
Database architecture, caching, and network latency	Infrastructure design choices can significantly affect whether an API performs consistently under load.
Authentication and authorization overhead	Security services are essential but also introduce processing steps that should be included in performance testing.

Consideration	Why It Matters
Rate limiting, error rates, and retry behavior	An endpoint should be evaluated not only for speed, but also for predictable behavior when requests fail or exceed expected thresholds.
Availability and sustained versus peak performance	Population health APIs must remain reliable over time, not only during short test windows or controlled demonstrations.
Source system and facade performance	An API facade can meet benchmarks while the source system behind it remains the constraint. Performance should be established at the source and at the facade, including laboratory, imaging, pharmacy, and EHR systems.

THE COST OF FULL POPULATION DIGITAL QUALITY MEASUREMENT

Moving to full population quality reporting is the right change to make because sampling cannot support population health management. FHIR is the only standards-based mechanism mature enough to make full population reporting operationally feasible. This shift, however, carries a cloud infrastructure cost that is unaccounted for in the current program design. Cloud computing for population-scale calculation is expensive, and vendor pricing for measure calculators varies widely with no reference point. The current model for cloud computing forces each ACO to negotiate its own pricing. ACOs are put at a disadvantage because current cloud computing discounts require guaranteed volumes and significant technical skills to optimize cloud costs for extraction and calculation (known as “egress” and “compute”).

Based on modeling the cost of full population reporting on reserved capacity rather than on-demand pricing, costs run **\$1.37–\$1.73** per member per year.³ The cost scales directly with population size and recurs annually regardless of program outcome. Three components drive the estimate. Storage covers retention of source data across the performance period. Compute covers four distinct stages: extraction, cross-source matching, normalization, and calculation. They are frequently modeled as a single step and understated as a result. Egress occurs twice, once for submission to CMS and separately to payers for Healthcare Effectiveness Data and Information Set (HEDIS) reporting.

³ Key assumptions behind the \$1.37–\$1.73/member range: reserved (not on-demand) cloud pricing throughout; quarterly processing cadence, standardized as the realistic floor given an ACO's core mission of managing total cost of care—neither CMS's minimum nor a monthly ACO preference; a 48–72 hour processing window; egress to CMS at four passes/year (quarterly) and to payers for HEDIS at one pass/year; compute modeled across four distinct stages (extraction, cross-source matching, normalization, calculation) rather than a single step; measure calculation anchored to a measure-by-measure classification of MSSP's measure concepts. Full derivation in Appendix 2.

Pricing for both cloud capacity and measure calculators is tiered by volume. The rates that make the low end of this range achievable are the rates cloud vendors and calculator vendors reserve for their largest customers. An ACO that reports on 15,000 covered lives does not reach those tiers and pays retail. The smallest reporting entities, therefore, carry the highest per-life cost, and the entities CMS most wants to bring into digital reporting are those that are most likely to have costs that exceed the value of participation.

CMS needs to offer policy guidance that will enable ACOs to work with major cloud vendors to support the transition to cloud-based computing. **Two options would relieve this situation without requiring CMS to store ACO clinical data.**

- **Facilitate shared infrastructure.** Computation could run through channels that already operate FHIR infrastructure at scale, including HIEs, the CMS Health Technology Ecosystem, and eventually the Trusted Exchange Framework and Common Agreement (TEFCA). Individual organizations would not provide standing capacity of their own, and the volume that qualifies for enterprise tier pricing would accumulate at the shared layer rather than for each participant.
 - This option employs existing federal authorities; however, open questions remain regarding whether HIEs have the bulk-FHIR maturity given that HIE technical capability varies by market. Whether quality measurement data flows fit TEFCA's eventual FHIR-based permitted purpose framework has yet to be resolved.
- **Facilitate volume pricing.** Federal cloud purchasing mechanisms could be extended to reporting entities as a group. Participants would buy at negotiated federal rates rather than at individually negotiated rates, which closes the gap for small entities without building new infrastructure.
 - This option is based on models used in other areas of federal policy but would probably require new legislation to create eligibility that has yet to be developed for any private-sector entity.
 - Through its health technology ecosystem, CMS could encourage cloud vendors to create a pathway for ACOs to use cloud computing for CMS quality reporting.

A NEW POLICY ROADMAP BUILT ON FHIR

To implement the proposed architecture, CMS should determine how interoperability capabilities are used within Medicare quality reporting programs while ONC continues to define and certify the interoperability capabilities required for standards-based exchange. In the [CY 2027 Physician Fee Schedule \(PFS\) proposed rule](#), CMS issued a request for information on a plan for moving to digital quality measurement supported by FHIR as well as other changes to the Medicare Shared Savings Program rules that show their desire to move to FHIR-based infrastructure. Still, regulatory changes, detailed below, are necessary to make this architecture possible without requiring it.

The policy shift is simple: CMS should define the required outcome—trusted, standards-based, auditable digital quality reporting—not require a specific technology category to produce it.

CMS can take a practical first step by pursuing the following actions:

- 1. Allow ACOs to report from any standards-based cloud FHIR endpoints instead of requiring reporting from EHRs.** The endpoints could aggregate clinical data, claims data, or both from participating providers and other permitted sources that the ONC recognizes as certified APIs as of July 2026. These endpoints come from trusted sources that aggregate data from participating providers and other permitted sources.
- 2. Use a common conformance framework.** Endpoints that ACOs use should demonstrate conformance through applicable ONC [Inferno](#) test scripts and other CMS-approved testing tools.
- 3. Make data quality explicit.** An open standard like the [PIQI Framework](#) can be used to determine whether data are complete, accurate, timely, and fit for a specific quality measurement use case.
- 4. Measure API performance at scale.** APIs that work for individual patient access should also perform across population-level queries.
- 5. Address the infrastructure costs directly.** CMS should either facilitate shared computation through already FHIR-capable channels or extend federal cloud purchasing to reporting entities collectively. Either approach closes the gap between individually negotiated and enterprise tier rates without requiring CMS to hold ACO clinical data.
- 6. Test before mandating.** CMS should start with a voluntary early adopter program involving ACOs, EHR vendors, technology companies, and measure developers to prove the approach in real-world settings.

To implement these recommendations, CMS should:

- **Redefine CEHRT at 42 CFR § 414.1305** to include any API certified under the ONC Health IT Certification Program to the standards adopted at 45 CFR 170.215.
- **Create a new Medicare digital quality measures (Medicare DQMs) data collection type as a voluntary pathway for quality reporting** that uses Standardized FHIR-based APIs in 42 CFR § 414.1305 and 42 CFR § 425.512. This change will allow ACOs to meet the quality performance standard if they report quality data using Medicare DQMs. Allowing pay for reporting for this option would be an incentive for organizations to use this option and would give CMS valuable data and the opportunity to work through implementation challenges without requiring organizations to put millions of dollars in shared savings at risk.
- **Add an option for meeting Medicare Shared Savings Program (MSSP) CEHRT use requirements at 42 CFR § 425.507 by reporting objective results for conformance, data quality, and performance of the FHIR APIs used for reporting** of at least one of the five ACO-reported measures in the Alternative Performance Pathway (APP) Plus measure set. ACOs would report “yes” for each measure by showing that the APIs passed conformance testing for the required FHIR standards through ONC’s Inferno Test Kit, met a data quality baseline such as the HL7 Patient Information Quality Improvement (PIQI) Framework, and met performance requirements for population-level queries. This option would tie the CEHRT requirement to the infrastructure ACOs use for quality reporting and would only be available to ACOs electing the Medicare DQM pathway.
- **Implement steps to relieve the cost of cloud computing necessary for Medicare DQMs**, by facilitating shared infrastructure or by finding ways to extend favorable volume pricing to entities undertaking Medicare DQM collection.

OTHER OBSTACLES TO DIGITAL QUALITY MEASUREMENT FOR ACOS

The framework proposed here addresses conformance, data quality, and performance for FHIR endpoints. It resolves only some of the challenges standing between ACOs and digital quality measurement. Several issues require separate attention from CMS and ONC.

- **Measure specifications call for data that the regulated APIs are not required to expose.** Specifications reference data elements that fall outside what a certified API must return, and they sometime point elsewhere than at the FHIR resources carrying the information. Imaging results arrive as PDFs under current attachment rules, which leaves ACOs performing manual work to make them usable for calculation. Specifications should align with the FHIR resources that reflect where providers document, such as the DiagnosticReport for breast cancer screening.
- **Much care documentation is unstructured.** Notes, attachments, and free text fields carry clinical content that no FHIR resource exposes discretely. Extraction reaches only what the source system has coded, and no standard addresses the remainder. Natural language processing is already used to address this problem, and other AI methods offer promise.
- **Real-time connectivity is new for ACOs and requires a different investment.** Establishing a connection includes networking work, IP whitelisting, authorization configuration, and service account provisioning across organizations, and it requires skilled resources that may be new for an ACO. Sustained connectivity also requires standing compute capacity, which is a recurring operating cost rather than a one-time build.
- **Measures are not aligned across programs.** A single clinical concept may appear in a dozen or more variants across the programs in which ACOs participate. ACOs that standardized years of custom EHR codes to HEDIS specifications find CMS requirements differ and that one code cannot satisfy both, which forces a choice between CMS reporting and supplemental data submission to commercial payers. Alignment of measures and value sets across reporting bodies would remove that decision-making burden and reduce costs of reporting.

CONCLUSION

ACOs are already responsible for managing populations and reporting quality results. The missing piece is regulatory alignment that allows those responsibilities to be supported by modern, standards-based infrastructure instead of approaches developed solely for quality measurement and reporting.

CMS has spent years establishing interoperability requirements; ONC has established FHIR certification requirements; Inferno provides an increasingly mature testing framework; HL7 is advancing standards for evaluating data quality through PIQI; and the private sector is building cloud-based data infrastructure capable of operating at population scale. CMS should allow qualified ACOs to use standards-based FHIR endpoints that aggregate data from participating providers, demonstrate conformance through Inferno, assess data quality through PIQI, meet measurable API performance requirements, calculate CMS-approved quality measures, and submit auditable results to CMS.

This approach preserves the EHR's essential role as a source of clinical information while allowing the ACO to use a population-level environment needed to aggregate, validate, calculate, and report quality information. It would also preserve federal oversight by giving CMS a more effective architecture to oversee—one that is standards-based, auditable, performance-tested, and capable of supporting multiple authorized uses.

ABOUT THE AUTHORS

**Ryan Howells***Principal, Leavitt Partners*ryan.howells@leavittpartners.com**Mark Marciante***Director, Leavitt Partners*mark.marciante@leavittpartners.com**Sarah Hudson Scholle***Principal, Leavitt Partners*sarah.scholle@leavittpartners.com**Daniela Simpson***Senior Consultant II, Wakely*daniela.simpson@wakely.com

ABOUT LEAVITT PARTNERS

Leavitt Partners, an HMA Company has been at the forefront of navigating change in health care for over a decade. We work at the intersection of government and health care, helping our clients stay one step ahead of the influences impacting the industry, including economic, market, delivery system, public policy, and political developments. Learn more at leavittpartners.com.

© 2026 Leavitt Partners. All Rights Reserved.

ABOUT WAKELY

Founded in 1999, Wakely Consulting Group, an HMA Company, is well known for its top-tier healthcare actuarial consulting services. With nine locations nationwide, Wakely boasts deep expertise in Medicare Advantage, Medicaid managed care, risk adjustment and rate setting, market analyses, forecasting, and strategy development. The firm's actuaries bring extensive experience across all sectors of the healthcare industry, collaborating with payers, providers, and government agencies.

© 2026 Wakely Consulting Group. All Rights Reserved.