

INFRASTRUCTURE
INTELLIGENCE
OPPORTUNITY

DETERMINISTIC BROADBAND

EXECUTIVE WHITE PAPER

Build the network for the market it will meet.

FIBER

RESILIENCE

SCALE

OPTIONALITY

A BRIGHTER
CONNECTED
TOMORROW

B=B
BIGBANG
BROADBAND LLC

Contents

Executive Summary.....	3
A 2026 Design Is Not a 2026 Network.....	3
The Market Is Already Telling Us the Direction.....	3
The Physical Decision Comes First.....	4
What ICQ Changes	4
The Economic Postulate	5
The Executive Decision.....	5
2. Determinism Comes First	6
The Requirement Begins at the Service Location	6
Determinism Begins With Availability	6
Why Determinism Precedes Capacity.....	6
From Requirement to Infrastructure Decision.....	7
3. The Market Will Change Before the Network Is Commissioned	7
The Deployment Clock Is Measured in Years	7
The Market Clock Is Measured in Months	8
The Change Is Broader Than Traffic Growth.....	8
The Planning Horizon Must Reach Commissioning.....	8
4. Capacity Is the Second Pressure, Not the First.....	9
Capacity Is No Longer a Single Number	9
Localization Moves Capacity Into the Local Network.....	10
Local Traffic Can Multiply Access-Side Work	10
Deterministic Behavior Defines the Capacity Requirement	10
5. Build the Network for the Market It Will Meet	11
A Responsive Network Is a Revenue Platform.....	11
The SLA Turns Capability Into a Product	11
Preserve the Option; Defer the Spend	12
Capability Is Also Competitive Defense	12
A Deterministic Network Creates Natural Partnerships	12
The Capital Test Is Lifetime Value.....	12
6. Conclusion: Do Not Commission Yesterday's Network	14
ICQ Turns Assumption Into Decision-Quality Evidence	14
From Qualification to a Prioritized Determinism Action Plan	14
The Business Benefits Compound Across the Ownership Life	15
Preserve the Capability Before the Opportunity Is Lost.....	15
Start With an Initial Discussion.....	16
Endnotes.....	17

THE CHEAPEST NETWORK TO BUILD MAY BE THE MOST EXPENSIVE NETWORK TO OWN

Build for today's demand, and tomorrow's revenue may require you to build again.

Executive Summary

This white paper advances a direct postulate: *a service provider that designs a multi-year network around only the demand visible when design begins is using the wrong planning horizon.* Further, if that horizon also drives economic decisions made in the name of fiscal efficiency, the design judged cheapest today is almost certainly not the design with the highest lifetime economic value. By the time that network is commissioned, it will be facing market demands that have continued to evolve throughout the entire design and construction cycle.

This is not a prognostication built on speculation. The market is already telling us where it is going. We cannot know the exact mix of applications, traffic, service classes, customer expectations, and revenue opportunities that will exist when the network is commissioned, but we do know that those conditions are changing faster than physical networks can be designed, permitted, built, and commissioned.

A 2026 Design Is Not a 2026 Network

For this paper, roughly four years from paper to full commissioning is used as a conservative stakeholder planning assumption for a major deployment. That is not a claim that every project takes exactly four years. It is based upon a schedule that can include design, rights-of-way, environmental and historic review, permitting, pole and make-ready coordination, procurement, supply-chain exposure, labor availability, construction, testing, acceptance, and rework. Any snag, such as a denied permit, changed route, environmental finding, or make-ready problem can send the project back to engineering and restart downstream steps.¹

That planning horizon is consistent with the realities reflected in current deployment programs and regulation: NTIA's BEAD framework requires funded projects to complete buildout and initiate service within four years after the subgrant; federal broadband permits on federal land can have 270-day processing deadlines; FCC large-order pole rules contemplate make-ready periods measured in months; and industry workforce studies continue to identify shortages in construction and technician labor. In other words, a network designed in 2026 may not be commissioned until 2030. Whatever assumptions shaped the design, that newly commissioned network will be facing the market demands of 2030—not merely those visible today.

The Market Is Already Telling Us the Direction

AI makes the speed mismatch impossible to ignore. In a matter of months, OpenAI has documented a rapid shift from short chatbot interactions toward delegated agentic work that can run for minutes or hours. By May 2026, 70.2 percent of sampled individual Codex users had made at least one request estimated to exceed an hour of human work.² Separately, OpenAI's real-time voice engineering identifies low and stable round-trip time, low jitter, and low packet loss as concrete requirements for natural interaction.³

Anthropic is moving in the same direction. Claude Sonnet 5 was introduced as an agentic model that can plan, use browsers and terminals, and operate autonomously; Claude Tag connects Claude to team channels, tools, data, and codebases and can pursue work asynchronously over hours or days. Anthropic is also

extending Claude into operational environments used by banks, airlines, insurers, manufacturers, government agencies, and professional-services firms.⁴

And AI is just a visible example, not the entire case. Edge and local computing, healthcare, public safety, industrial automation, video and security, enterprise platforms, utilities, interactive media, and other workloads are also changing where applications run, where traffic originates and terminates, and the value customers place on continuity, traffic priority, predictable performance, and rapid recovery. ETSI and AWS already describe edge-computing environments designed specifically to place compute closer to users for low-latency and resilient applications.⁵

Network planners have always had to design beyond the current traffic snapshot. For decades, that discipline often meant planning for relentless traffic growth. The present challenge is broader: not merely more traffic, but changing traffic patterns, application locations, service expectations, and revenue opportunities. Consequently, the planner is no longer dealing only with growth in quantity. The planner is dealing with a fundamental change in the very nature of the communications the network is tasked to support.

Where deep carrier planning experience and expertise exist, that discipline may already be institutionalized. Where it does not, the design method itself must force the right questions: What will this network need to be capable of when it actually enters service? Which emerging requirements are already visible? Have we identified which physical options must be preserved now so the network can respond as those requirements mature? And if additional capability is needed later, can it be added through equipment, splicing, activation, and configuration—or will the provider have to construct again?

The Physical Decision Comes First

The physical network is the slowest part of the system to change. Electronics can be installed. Fibers can be spliced. Capacity can be lit. Network-level service differentiation and traffic priority can be configured. Protection can be activated. But those actions work only if the underlying physical options exist: routes, conduit, fiber, space, access, diversity, alternate entrances, and protection relationships.

If an alternate route was never preserved, software cannot create it. If reserve fiber does not exist, it cannot be lit. If a backup path shares a single point of failure with the primary path, it may be redundant, but it is plainly not diverse.

That distinction is the economic hinge. A turn-up adds equipment, splices, configuration, and activation to physical infrastructure already capable of supporting the service. Reconstruction starts over with permits, engineering, crews, excavation, remobilization, route work, entrances, or cable replacement—while the revenue opportunity waits.

What ICQ Changes

ICQ is not a prescription to build every imaginable capability now. It is a method for deciding which physical choices must be provided or preserved during the current design and construction cycle so the network can respond to the market it enters—and continue adapting throughout its ownership life—without repeatedly returning to construction.⁶

That may mean sufficient conduit, fiber reserve, a preserved opportunity for an independent route, adequate space and access, alternate entrances, or avoidance of common points of failure. The electronics, optical equipment, splices, configurations, protection schemes, and services can wait until demand and revenue justify them. An expensive-to-recover physical option should not be eliminated simply because the service that may eventually use it has not yet been activated.

That is capital discipline, not overbuilding. ICQ preserves economically valuable choices without pretending management can predict exactly which ones will be monetized, when, or at what scale. It makes visible which choices are practical to preserve during construction and which become expensive to recover after construction.

The objective is not to spend more. It is to avoid spending twice—and to avoid commissioning a new network that is already physically constrained against a market that moved while the network was being built.

The Economic Postulate

A provider that preserves the right physical options can defer much of the service-specific spending until demand appears while retaining the ability to act quickly when it does. A provider that removes those options may save capital at first but then incur greater lifetime cost through reconstruction, delayed revenue, or services it cannot economically offer.

That is why this paper advances the postulate that the design judged cheapest today is almost certainly not the design with the highest lifetime economic value. The burden of the paper is to test that claim against market evidence, deployment realities, network architecture, and the economics of preserving versus reconstructing capability.

The Executive Decision

When the market demands a service the network was not originally activated to provide, will the answer be equipment, splices, configuration, and activation—or another construction project?

2. Determinism Comes First

Capacity matters, but it is not the first question. Before a provider asks how much traffic a network can carry, it should understand what the network must be able to do for the locations that depend on it. That is the starting point for determinism.

In this paper, determinism is not treated as a single network-wide feature or a universal standard. It is relative to the service location and to the work being performed there. A residence, hospital, public-safety answering point, enterprise, utility facility, wireless site, school, AI compute location, or industrial operation can place very different demands on the same broadband network. The appropriate level of availability, protection, restoration, latency, capacity, traffic priority, and resilience follows from the function of the location and the consequences of service degradation or failure.

The Requirement Begins at the Service Location

The practical question is not, “Is this a deterministic network?” The more useful question is, “What must this network be able to do for this location?” For one location, ordinary best-effort broadband may be entirely appropriate. For another, even a short interruption may disrupt emergency communications, clinical operations, industrial processes, security systems, financial activity, or another workload whose value depends on continuity and predictable performance.

That distinction is why Big Bang Broadband begins with Service Location Information Collection (SLIC). SLIC develops and governs the information ICQ needs about each location: what the location does, what workloads it supports, how critical those workloads are, and what communications requirements follow. The governed SLIC information includes availability, protection and restoration needs, capacity, latency, security sensitivity, and optical optionality, among other attributes.⁷

The sequence is deliberate: the service location defines the requirement; SLIC captures and governs that requirement; ICQ determines whether the physical and optical infrastructure can support it. SLIC does not itself qualify the network. ICQ does not invent the service requirement. The two processes are linked because infrastructure cannot be judged adequate in the abstract; it must be judged against what the network is actually expected to support.

Determinism Begins With Availability

Determinism also begins with a condition that can be overlooked when the discussion focuses too quickly on software, traffic management, service differentiation, or traffic priority: **the network has to be there when it is needed.** A service cannot be predictably delivered across a path that disappears when a cable is cut, a structure fails, a common conduit is damaged, or a shared point of failure removes both the primary and the nominal backup connection.

Reliability, redundancy, route diversity, protection paths, restoration capability, and independence from common failure points are therefore not separate from determinism. They are part of its physical foundation. Network-level service differentiation and traffic priority become meaningful only after the network has viable paths and resources across which those controls can operate.

This is the point at which physical design decisions become economically consequential. Software can choose among routes only if alternate routes exist. Protection can switch traffic only if a protection path has been provided. Additional fiber can be activated only if fiber, conduit capacity, access, splice opportunities, and other required physical conditions were preserved. A configuration change can use an option; it cannot create a physical option that construction eliminated.

Why Determinism Precedes Capacity

A capacity-first discussion can obscure this distinction. A network may have substantial unused bandwidth and still be unable to support a location whose requirement includes higher availability, independent protection, faster restoration, or differentiated treatment. More bandwidth does not create a second route. It

does not remove a shared bridge crossing. It does not create an independent building entrance. And it does not determine which traffic should receive priority when multiple services compete for network resources.

Determinism can also improve the economic use of capacity already in place. Where the network and its operating systems permit it, a provider can differentiate services, prioritize important traffic, reserve resources, apply protection selectively, and align network behavior with the value and criticality of the workload. That does not eliminate the need for additional capacity as demand grows. It means that capacity should be considered after the provider understands the service behavior it is trying to deliver.

This is also why a single protection standard imposed everywhere would miss the point. ICQ is not a doctrine of maximum redundancy or maximum expenditure. It is a qualification method that asks whether the infrastructure serving each location contains the physical and optical capabilities appropriate to that location's requirements. The SLIC process exists so those requirements are evidence-based rather than assumed.

From Requirement to Infrastructure Decision

The distinction becomes especially important while a network is still being designed. Some capabilities can be deferred economically. Electronics can be added later. Fibers can be spliced. Capacity can be lit. Protection can be activated. Network-level service differentiation and traffic priority can be configured when the service and revenue opportunity justify them. But those later actions depend on physical choices made much earlier.

The objective is therefore not to equip every location for the highest imaginable requirement. It is to identify where the physical network must preserve choices so that a higher service requirement can be satisfied without another construction program. That is the relationship between SLIC and ICQ: location-specific requirements are developed first, and the infrastructure is then qualified against them before the opportunity to preserve those choices is lost.

This relationship is already being carried into engineering work. The current Biarri Networks-BBB program is structured around service-location protection classifications and uses those requirements to evaluate best-effort outside-plant designs and identify protection shortfalls.⁸ That work is moving the determinism discussion from principle to measurable infrastructure requirements.

3. The Market Will Change Before the Network Is Commissioned

Network planning is governed by two clocks that move at very different speeds. The physical network moves from design to full commissioning over years. The applications, computing architectures, customer expectations, and service opportunities it must support can change materially in months. A coherent planning horizon has to account for both. Otherwise, the network can be technically faithful to the design package and still enter service behind the market.

The Deployment Clock Is Measured in Years

The roughly four-year paper-to-commissioning assumption used in this paper is not based on the idea that construction itself necessarily takes four years. It reflects the sequence of work that surrounds construction: engineering, rights-of-way, environmental and historic review, permitting, pole and make-ready coordination, procurement, labor availability, construction, testing, acceptance, and the possibility of rework. Some activities can overlap. Others cannot. More importantly, a problem discovered downstream can force the project back upstream. A denied permit, an environmental finding, an unavailable route, or a make-ready problem can require redesign and can restart approvals, material decisions, scheduling, and construction dependencies.⁹

Current deployment evidence reinforces that this is not merely a theoretical project-management concern. In the Fiber Broadband Association and Cartesian 2025 deployment-cost survey, 46 percent of respondents reported material project delays. Permitting, utility locates, and make-ready work were the most frequently cited causes. Looking into 2026, 62 percent expected deployment timelines to be longer than normal. The same report notes that delay can create financing and opportunity costs because expenses continue while

revenue generation is postponed.¹⁰ A multi-year planning horizon is therefore not an argument about an exact completion date. It is recognition that a major network is exposed to enough sequential dependencies and rework risk that the market can move substantially before the asset is fully commissioned.

The Market Clock Is Measured in Months

While the deployment clock advances in years, the application clock can move in months. OpenAI reported that by May 2026, 70.2 percent of sampled individual Codex users had made at least one request estimated to exceed an hour of human work, and 25.6 percent had made at least one request estimated to exceed eight hours. The significance is not the specific percentages. It is the rapid movement from short, human-paced interactions toward software that can continue working independently, invoke tools, interact with other systems, and operate in parallel.¹¹

The service characteristics are changing as well. OpenAI's real-time voice engineering identifies low and stable round-trip time, low jitter, and low packet loss as concrete requirements for natural interaction.¹² Anthropic is advancing the same broader shift through agentic models that plan, use browsers and terminals, connect to enterprise tools and data, and pursue delegated work asynchronously over extended periods.¹³ These developments do not, by themselves, prove that every broadband service requires deterministic treatment. They do demonstrate how quickly application behavior and network sensitivity can change during a single infrastructure-development cycle.

Netflix provides a mature example that localization is not a speculative AI-era concept. Through Open Connect, Netflix says it partners with more than a thousand ISPs to localize substantial amounts of traffic using appliances embedded within ISP networks; each embedded deployment can offload substantial content traffic from peering or transport circuits, and multiple deployments can be distributed geographically or across the network to maximize local offload.¹⁴ The point is not Netflix itself. It is the architecture: major application traffic can be brought inside the provider network, changing where traffic is handed off and carried. The current edge-computing movement extends that localization model from cached content toward compute, inference, and operational applications.

The same direction is visible in the computing infrastructure being built around those applications. AWS places compute and storage inside communications-service-provider networks through Wavelength.¹⁵ Microsoft places Azure Extended Zones in metropolitan areas and industry centers for low-latency and data-residency workloads. Google Distributed Cloud extends cloud and AI infrastructure into data centers and edge locations for local processing, survivability, low-latency operation, and edge inference. AWS identifies use cases that include healthcare, real-time video analytics, connected vehicles, smart factories, and other applications in which processing is moved closer to where data is generated or consumed.¹⁶

The Change Is Broader Than Traffic Growth

Network planners have always designed beyond the current traffic snapshot. Historically, much of that discipline could be expressed as growth: more subscribers, higher peak demand, larger upstream and downstream loads, and the capacity expansion required to stay ahead of them. That discipline remains necessary, but it is no longer sufficient. The network is being asked to support changes in where applications run, where traffic originates and terminates, how long machine-driven activity can continue without human pacing, how sensitive some applications are to delay and loss, and how consequential a service interruption can become.

Those changes do not all point to the same network requirement, and they will not mature uniformly across every service territory. That is precisely why the planning method matters. The evidence is not sufficient to tell a provider exactly which application or service class will dominate its market at commissioning. It is more than sufficient to reject the premise that the demand visible when design begins is an adequate baseline for the network that will emerge from a multi-year build.

The Planning Horizon Must Reach Commissioning

The planning problem is therefore not an inability to predict the future with precision. *The uncertainty is in the destination, not the direction.* The direction is already visible in the market: applications are becoming more

distributed, more integrated with operational systems, more capable of acting without continuous human input, and in many cases more sensitive to locality, continuity, and predictable network behavior. A service provider does not need to know the exact 2030 application mix to know that a 2026 demand snapshot will not describe it.

The appropriate response is not to equip every possible service in advance. It is to evaluate design decisions against the market the network will meet when it is commissioned and to avoid eliminating physical choices that visible market direction already suggests matter. The immediate consequence is that capacity can no longer be treated as the singular measure of how much more of today's traffic the network will carry. Application location and traffic behavior are changing as well. More importantly, once deterministic network behavior becomes a service requirement, capacity is no longer the primary design consideration. It becomes a secondary, service-dependent question: whether sufficient network resources exist, in the right places and under the right operating conditions, to deliver the behavior the application requires.

4. Capacity Is the Second Pressure, Not the First

Section 2 established the governing sequence: determine what the network must do for the service location, then determine whether the network has the resources to do it. Once deterministic behavior enters the requirement, capacity stops being a single bulk measure of throughput. It becomes service-dependent.

Capacity Is No Longer a Single Number

Traditional capacity planning can be expressed largely as a sizing problem: how much traffic the network carries, how fast that traffic grows, and how much reserve is needed to stay ahead of demand. Once a service requires deterministic network behavior, that is no longer enough. The provider must ask whether sufficient capacity exists at the required location, in the required direction, on the required path, and under the operating or failure condition in which the service must continue.

A standards-based example makes that distinction explicit. The IETF Deterministic Networking architecture describes deterministic service as the dedication and allocation of network resources such as link bandwidth and buffer space to specific flows or classes of flows along their paths. It identifies resource allocation, service protection, and explicit routes as the principal techniques used to deliver deterministic qualities of service, and states that bounded delivery latency is achieved by reserving bandwidth and buffer resources at each DetNet node along the flow path. The relevant measure is therefore not simply aggregate network capacity; it is provisioned capacity and other resources on the path that must deliver the required behavior.¹⁷

A network can have abundant aggregate bandwidth and still be unable to satisfy a service requirement. Protected traffic is of little value if the alternate path lacks the capacity to carry it after a failure. Downstream headroom does not solve an upstream constraint. Capacity available elsewhere in the network does not relieve a bottleneck on the feeder, access segment, aggregation path, or interconnection point that the workload actually uses.

That changes the economic role of capacity. The objective is not maximum capacity everywhere. It is sufficient usable capacity aligned with workload criticality and required network behavior. For ordinary best-effort traffic, available capacity can be shared opportunistically. For a service that must continue through a failure or receive differentiated treatment during congestion, the provider must ensure that the necessary resources remain available when the required condition actually occurs. Capacity therefore becomes relative to service class, direction, location, path, and failure state.

The same principle appears in the access standard itself. ITU-T G.9807.1 treats quality of service as a resource-allocation problem for individual traffic flows and traffic-flow aggregates, provisions flows with distinct upstream and downstream service parameters, and assigns QoS-aware traffic-management responsibilities to the OLT and ONU based on service specifications, available resources, bandwidth availability, and dynamic traffic conditions. Its upstream reference model distinguishes fixed, assured, non-assured, and best-effort bandwidth components. A nominal 10-Gbit/s symmetric PON is therefore not

engineered as a single undifferentiated capacity pool when service behavior matters; capacity is scheduled and allocated according to service requirements and direction.¹⁸

Localization Moves Capacity Into the Local Network

Section 3 established that localization is already a production architecture, not a speculative AI-era concept. Netflix Open Connect brings substantial content traffic inside ISP networks,¹⁹ while AWS, Microsoft, and Google are extending compute, inference, and application infrastructure toward metropolitan, provider, and edge environments.²⁰ A provider should not interpret that movement as traffic disappearing. Localization changes where traffic is carried.

When an application is remote, the subscriber is typically the local endpoint and the other endpoint sits beyond the local access network. When the application itself moves inside the provider footprint, both endpoints can become local. The provider can then carry the subscriber-side access leg and the application-side access leg of the same useful transaction. Long-haul or upstream transport can fall while access, aggregation, metro, and interconnection loading closer to the endpoints increases.

Local Traffic Can Multiply Access-Side Work

A simplified PON example makes the point. With a remote application, a subscriber request crosses the subscriber PON upstream and the response returns downstream. If the application endpoint is instead attached to another ONU on the same PON, the request travels upstream from the subscriber and downstream to the application endpoint; the response then travels upstream from the application endpoint and downstream to the subscriber. Ignoring protocol overhead, the application payload therefore consumes roughly twice the aggregate PON-facing transport for the same useful transaction.²¹

The directional effect can be even more significant for asymmetric workloads. A small request followed by a large response is strongly downstream in the remote-application model. When both endpoints are local on the same PON, both the request and the response contribute traffic in both directions, making the access load far more symmetric. The exact load placed on an OLT fabric, uplink, or aggregation network depends on topology and where forwarding or routing occurs, but the planning point remains: moving an application closer does not automatically reduce the capacity required inside the local network.

Deterministic Behavior Defines the Capacity Requirement

The same principle applies to protection and service differentiation. If a service must survive a feeder or path failure, the surviving path must have enough capacity to carry the protected workload under that failure condition. If a service must receive priority when resources are contested, sufficient schedulable capacity and headroom must exist where that contention occurs. If the service requirement includes predictable latency or loss performance, the network must have enough usable capacity, together with the appropriate traffic controls, to sustain that behavior under the conditions in which the commitment matters.

Protection makes failure-state capacity concrete. IETF RFC 3469 distinguishes an equivalent recovery path that can replace the working path without degrading service from a limited recovery path that lacks the resources or resource reservations to do so. In 1+1 protection, recovery-path bandwidth, buffers, and processing capacity are fully reserved and carry the protected traffic. In 1:1 protection, recovery-path resources can carry preemptible lower-priority traffic until a fault occurs, at which point the protected traffic takes those resources. Protection capacity therefore has meaning only in relation to the failure state the network is designed to survive.²²

Commercial service standards express the same principle from the service side. MEF 61 associates classes of service with defined performance objectives and related bandwidth profiles for IP services. MEF 10.4 applies bandwidth profiles to ingress and egress Class of Service flows and defines envelopes in which ranked bandwidth-profile flows can share and reallocate unused bandwidth. The service model therefore ties bandwidth to service class and required behavior rather than treating total interface rate as the only capacity measure.²³

Capacity has not become less important. It has been repositioned. Deterministic behavior defines what the network must do; capacity is one of the resources required to do it. Nominal headroom on the wrong path, in the wrong direction, at the wrong location, or unavailable during the required failure state does not satisfy the service requirement simply because that capacity exists somewhere in the network.

The planning question is therefore no longer only how much aggregate capacity to build. It is what capacity must exist, where it must exist, and under what operating conditions it must remain available to support the requirements defined at the service location. Once capacity is understood that way, the next question becomes physical and economic: which network options must be built or preserved during deployment so those resources and behaviors can be activated without another construction program?

5. Build the Network for the Market It Will Meet

Sections 2 through 4 established the technical premise: determine what the network must do for the service location, then determine whether the network has the resources to do it. The business consequence is more direct. A broadband network is a long-lived revenue asset entering a market that will continue to change throughout its ownership life. Its economic value depends not only on how cheaply it can be built, but on how many commercially useful choices remain available after it is built.

The market's increasing sophistication is not a planning burden. It is an opportunity to differentiate. A provider whose network can support different levels of availability, protection, restoration, latency, traffic treatment, and locality can offer more than another speed tier. It can match network behavior to customer value. That direction is already visible in the edge market: AWS describes Wavelength as infrastructure for workloads that require low latency or edge resiliency inside communications-service-provider networks, while GSMA's Operator Platform work explicitly contemplates exposing capabilities such as quality of service and low-latency interconnection to enterprise customers and application developers so operators can monetize those capabilities.²⁴

A Responsive Network Is a Revenue Platform

Where providers sell substantially the same best-effort product, competition gravitates toward speed, price, promotion, and bundle. A responsive network changes the commercial surface. The provider can create offers around the behavior that different customers actually value: continuity for a critical site, protection for a mobile or enterprise facility, lower and more predictable latency for an edge workload, faster restoration for an operational system, or differentiated treatment for traffic whose interruption carries greater consequence.

That is the economic meaning of sustainability in this paper. A sustainable network is one that can continue producing relevant, profitable service offers as the market changes. Network capability does not create profit by itself; it preserves the commercial choices from which profit is built. The more of those choices the physical network preserves, the less the provider is forced to compete only on the commodity attributes that every other network can advertise.

The SLA Turns Capability Into a Product

Sophisticated buyers do not have to accept a generalized assurance that a network is reliable or high performance. The provider can put the requirement in writing. MEF's work on shared 5G transport describes deterministic service-level specifications that can be enforced for bandwidth, low latency, and high availability, and defines the Service Level Specification as the technical portion of a Service Level Agreement, including service performance, objectives, metric definitions, measurement, classification, and bandwidth-profile details.²⁵

That changes the revenue proposition. An SLA converts engineering sophistication into a contractual product. The customer is not being asked to trust a marketing claim; the provider is accepting measurable service obligations. That is a well-earned opportunity for additional revenue because the provider is selling behavior that ordinary best-effort service does not promise. As edge compute, AI, industrial, healthcare, public-safety, and other operational applications become more consequential, the ability to demonstrate and contract for network behavior becomes part of the product itself.

Preserve the Option; Defer the Spend

None of this requires every advanced service to be fully activated on day one. The sophistication can be turned up when a customer, application, partnership, or revenue opportunity requires it. What cannot wait is the physical opportunity to create that capability later. ICQ is designed to make that distinction visible by qualifying the physical and optical infrastructure, identifying capability gaps, and preserving evidence about protection, resilience, diversity, reserve, future optical capability, and other choices that become expensive to recover after construction.²⁶

That is capital discipline. Electronics, splicing, service logic, protection activation, and other service-specific spending can follow demand. Routes, conduit, fiber reserve, diverse entrances, access, space, and protection relationships cannot be assumed into existence after the fact. The ability to build the sophisticated service later has to survive the decisions made now. In economic terms, the provider preserves the option before it knows exactly when the option will be exercised.

Capability Is Also Competitive Defense

Underbuilding also creates competitive exposure. A provider that builds only enough to satisfy today's visible demand can leave an opening for a competitor to overbuild the market with a more capable network and a stronger service proposition. That is not a theoretical market structure. CostQuest's January 2026 analysis of FCC broadband data reports that approximately 28 percent of new fiber was being built to homes already served by both fiber and cable, and that smaller providers were directing roughly 40 percent of their new builds toward second-fiber markets. The same analysis identifies customer losses for cable operators in newly competitive fiber geographies.²⁷

Those figures do not say that every service territory will be overbuilt. They show that overbuilding is an active competitive strategy. If the existing provider cannot answer a more capable entrant without another construction cycle, the disadvantage is both financial and temporal: capital has to be remobilized while the competitor is already selling. Preserving capability does not eliminate competition. It reduces the opportunity for a competitor to define the market around capabilities the existing network physically cannot provide.

A Deterministic Network Creates Natural Partnerships

The same capability that differentiates retail and enterprise services also expands the wholesale and partnership opportunity. Cellular providers are a natural example. MEF's 5G transport work notes that mobile operators without their own transport facilities use services from local transport providers, and it identifies new revenue opportunities for service providers supplying mobile fronthaul and backhaul. It also treats mobile transport as a service with defined requirements for latency, bandwidth, synchronization, monitoring, and service-level performance.²⁸

A deterministic network makes that partnership proposition easier to evaluate. The cellular provider does not have to infer capability from a fiber map or accept an assurance that the network is robust. The infrastructure can be qualified against defined requirements, the service can be engineered to those requirements, and the obligations can be written into an SLA. The same logic extends to edge-compute operators, application developers, cloud providers, utilities, healthcare systems, public-safety organizations, and enterprises whose own products or operations depend on predictable communications behavior.

The Capital Test Is Lifetime Value

The financial question is therefore larger than first cost. It is the relationship among the cost to preserve capability during the original build, the cost to recover it later, the revenue that becomes available when the capability can be activated, and the time required to respond when the market asks for it. That is the difference between a network designed as a finished product and a network designed as a platform.

The current Biarri Networks-BBB protection-analysis program is designed to put part of that question into measurable terms. Its Type-B phase compares three cases for the same representative network: a best-effort design, that design corrected after the fact to satisfy the protection requirement, and a Type-B-optimized design built to the requirement from the start. The program will produce a three-way bill-of-materials and cost

comparison and isolate the retrofit premium. Results are not yet available; the purpose of the study is to create the evidence base rather than assume the answer.²⁹

That is the correct executive test. The question is not how cheaply a network can be made to carry today's traffic. It is what it costs to preserve the right to support higher-value services, partnerships, and contractual performance in the market the network will actually meet—and what it costs to recover that right after construction has already eliminated it.

Building for the market the network will meet does not require management to predict the exact winning applications of 2030. It requires management to avoid letting 2026 design decisions make 2030 commercial decisions in advance. A responsive network can differentiate, support enforceable SLAs, create partnership revenue, respond to capability-based overbuild, and activate higher-value services as demand appears. The objective is not maximum capability on day one. It is to preserve commercially valuable optionality at a rational capital cost.

6. Conclusion: Do Not Commission Yesterday's Network

The postulate advanced at the beginning of this paper has now been tested against the evidence, and the evidence proves it correct. The deployment realities, market changes, application requirements, deterministic service demands, traffic behaviors, protection requirements, capital consequences, and optionality considerations documented throughout this paper—and supported by the accompanying endnotes—lead to the same conclusion: a service provider that designs a multi-year network around only the demand visible when design begins is using the wrong planning horizon, and the design judged cheapest at the outset is almost certainly not the design with the highest lifetime economic value.

The question is therefore no longer whether the market will change before the network is fully commissioned. The question is whether the network will be capable of responding when it does.

ICQ Turns Assumption Into Decision-Quality Evidence

Infrastructure Capability Qualification changes the executive discussion because it replaces assumption with evidence. ICQ can show what the physical and optical infrastructure is actually capable of supporting against the requirements that matter at the service location; where availability, protection, diversity, restoration, capacity, reserve, or optical optionality are sufficient; where capability gaps exist; where a claimed backup relationship still shares a common failure point; where future choices have been preserved; and where missing or unresolved information must be closed before management relies on the result. The output is not another generalized statement that the network is robust, fiber-rich, or future-ready. It is a traceable capability picture tied to the network, the requirement, and the evidence.³⁰

That visibility creates business value before a single remediation project is approved. Management can distinguish between capability that already exists, capability that can be activated through electronics, splicing, configuration, or service logic, and capability that requires a physical intervention. It can see where the network is already positioned to support differentiated services and where design choices are quietly eliminating revenue options. It can also identify unsupported assumptions before those assumptions become capital commitments, contractual promises, or construction that has to be revisited later.

From Qualification to a Prioritized Determinism Action Plan

ICQ should not leave the owner with a flat list of deficiencies. Its candidate process is designed to identify evidence-supported alternatives, qualify technical sufficiency first, and then compare the consequences of the alternatives that remain. The comparison can address requirement fit, evidence completeness, route and failure-domain separation, constructability, reserve and optionality preservation, required additions, future rework, and capital-efficiency evidence where that evidence is available. The resulting candidate recommendations give the customer a disciplined basis for turning the qualification into a prioritized determinism action plan rather than treating every issue as equally urgent or equally valuable.³¹

The practical benefit is prioritization based on evidence and quantifiable benefit. Where the controlled technical, quantity, and cost evidence supports the comparison, management can determine what must change, why it must change, which service locations or workloads benefit, which deterministic requirement is satisfied, what physical work is required, what later rework is avoided, what reserve or optionality is preserved, and what bill-of-materials or capital consequence distinguishes one technically qualified alternative from another. The result is a clearer answer to the capital question: where does the next dollar of infrastructure spending produce the greatest supported improvement in capability and commercial flexibility?

ICQ does not take the investment decision away from the owner, and it does not convert a technically attractive candidate into an approved engineering design. It establishes the qualified decision population and the evidence supporting it; the authorized client or project authority chooses the business and capital direction; and the authorized engineering process converts that accepted direction into design. That separation is important because it keeps technical sufficiency ahead of cost preference while preserving management authority over capital, timing, risk, and commercial strategy.³²

The Business Benefits Compound Across the Ownership Life

The economic benefit is larger than avoided reconstruction. ICQ gives the provider a way to preserve choices that can improve lifetime return from the network: defer service-specific spending until demand exists; turn up higher-value capability without another civil program; reduce time to revenue when a new application or customer requirement appears; support differentiated service levels instead of competing only on speed and price; create a stronger basis for enforceable SLAs; improve the credibility of enterprise, cellular, edge-compute, AI, utility, healthcare, public-safety, and other partnership discussions; preserve the ability to respond when a competitor overbuilds with a more capable platform; and reduce the risk that a profitable opportunity has to wait behind permitting, engineering, excavation, make-ready, or route reconstruction.

ICQ can also reveal opportunity, not just deficiency. A qualified protection relationship can support a service class that ordinary best-effort infrastructure cannot credibly promise. Preserved route diversity can strengthen a mobile-transport or enterprise proposition. Reserve fiber, conduit, space, access, alternate entrances, and optical flexibility can become monetizable options when the market creates the demand. Conversely, a capability gap identified before construction can be addressed while the design is still fluid instead of becoming an expensive constraint after the network is in the ground. In each case, the value is the same: the provider retains the ability to respond.

That is the economic meaning of sustainability. A sustainable broadband business is not one that merely finishes construction and begins billing customers. It is one whose infrastructure continues to support relevant, differentiated, profitable services as the market changes. Preserving capability improves capital efficiency because it avoids both forms of waste: premature overbuilding of service-specific equipment that the market has not yet justified, and underbuilding of physical infrastructure that later forces the provider to construct again.

Preserve the Capability Before the Opportunity Is Lost

The network does not need to be fully equipped for every advanced service on day one. The capability can be activated when the customer, application, partnership, or revenue opportunity justifies the spend. What cannot be deferred indefinitely is the physical opportunity to support that capability. Routes, pathways, reserve fiber, diversity, access, space, protection relationships, and other long-lived choices become expensive—or impossible—to recreate after construction.

That is why ICQ belongs early enough in the planning process to influence the choices that are still recoverable. Build the physical option once. Turn up the capability when demand requires it. The objective is not to predict every service the market will demand. It is to make certain the network preserves the ability to respond when that demand arrives.

Start With an Initial Discussion

If you are planning a new broadband network, evaluating an engineering design, or deciding how to strengthen an existing network plan before capital and construction decisions become difficult to reverse, I invite you to contact me directly. An initial discussion can establish where ICQ fits within your network plans, what questions the qualification should answer, what service-location information is required, and how the resulting evidence can support a prioritized determinism and capital action plan.

David J. Malfara, Sr.
Chief Executive Officer
Big Bang Broadband LLC
Email: dmalfara@bigbangbroadband.com
www.bigbangbroadband.com

Endnotes

1. The roughly four-year paper-to-commissioning horizon used here is a conservative stakeholder planning assumption, not a claimed industry-wide median. Supporting context includes: NTIA, "BEAD Initial Proposal Guidance," requiring complete buildout and initiation of service within four years after a subgrant; U.S. Government Accountability Office, "Broadband Deployment: Agencies Should Take Steps to Better Meet Deadline for Processing Permits" (Apr. 10, 2024), describing a 270-day federal-land permit-processing deadline and missed deadlines; FCC 25-38, showing large-order pole make-ready periods extending to 120 days in communications space and 180 days above communications space after payment, while identifying permitting delays, workforce shortages, contractor scarcity, and material availability as schedule factors; and Fiber Broadband Association / Power & Communication Contractors Association, "Broadband Market Workforce Needs" (2024), estimating an immediate need for 28,000 additional broadband construction workers and 30,000 technicians. [NTIA source](#) [GAO source](#) [FCC source](#) [FBA/PCCA source](#)
2. OpenAI, "How agents are transforming work," June 25, 2026. OpenAI reports that agentic AI shifts knowledge work from short interactions toward delegated, long-horizon tasks and that, by May 2026, 70.2% of sampled individual Codex users had made at least one request estimated to exceed one hour of human work. [Source](#)
3. OpenAI, "How OpenAI delivers low-latency voice AI at scale," May 4, 2026, identifying low and stable media round-trip time, low jitter, and low packet loss as concrete requirements for natural real-time voice interaction; and OpenAI, "Build more natural voice experiences with GPT-Live-1 in the API," Sept. 10, 2026, extending full-duplex voice capability into developer applications and business workflows. [Engineering source](#) [API source](#)
4. Anthropic, "Introducing Claude Sonnet 5," June 30, 2026, describing planning, browser and terminal use, and autonomous operation; Anthropic, "Introducing Claude Tag," June 23, 2026, describing access to selected channels, tools, data and codebases and asynchronous work over hours or days; and Anthropic, "DXC will integrate Claude into the systems banks, airlines, and other regulated industries rely on," June 11, 2026. [Sonnet 5 source](#) [Claude Tag source](#) [DXC source](#)
5. ETSI, "Multi-access Edge Computing (MEC)," describing cloud-computing capabilities and an IT service environment at the network edge, including fixed access, with ultra-low latency and high bandwidth; and Amazon Web Services, "AWS Wavelength," describing compute and storage in telecommunications-provider data centers for low-latency, resiliency, and edge applications, including ML inference. [ETSI source](#) [AWS source](#)
6. Big Bang Broadband LLC, "ICQ - Client Executive Overview," Sept. 12, 2026. The source describes ICQ as beginning with the engineering-approved design and governed service-location information; qualifying the physical and optical infrastructure for purpose-specific capability; preserving unresolved conditions rather than assumptions; identifying capability gaps; and supporting decisions on protection, resilience, diversity, capacity, reuse, reserve, shared use, future optical capability, emerging technologies, capital allocation, and preservation of future choices.
7. Big Bang Broadband LLC, "ICQ SLIC Client Process Overview" (September 2026), communications/workload requirements and location-specific engineering inputs.
8. Biarri Networks and Big Bang Broadband LLC, "Type-B and Dual-Parent Protection Analysis Program" SOW (2026), Phases 1-4.
9. The roughly four-year paper-to-commissioning horizon used here is a conservative stakeholder planning assumption, not a claimed industry-wide median. Supporting context includes: NTIA, "BEAD Initial Proposal Guidance," requiring complete buildout and initiation of service within four years after a subgrant; U.S. Government Accountability Office, "Broadband Deployment: Agencies Should Take Steps to Better Meet Deadline for Processing Permits" (Apr. 10, 2024), describing a 270-day federal-land permit-processing

deadline and missed deadlines; FCC 25-38, showing large-order pole make-ready periods extending to 120 days in communications space and 180 days above communications space after payment, while identifying permitting delays, workforce shortages, contractor scarcity, and material availability as schedule factors; and Fiber Broadband Association / Power & Communication Contractors Association, "Broadband Market Workforce Needs" (2024), estimating an immediate need for 28,000 additional broadband construction workers and 30,000 technicians. [NTIA source](#) [GAO source](#) [FCC source](#) [FBA/PCCA source](#)

10. Fiber Broadband Association and Cartesian, "Fiber Deployment Cost Annual Report 2025" (Jan. 2026). The survey reports that 46% of respondents faced material project delays in 2025; permitting, utility locates, and make-ready work were the most cited delay drivers; and 62% expected deployment delays in 2026. The report also notes that extended timelines can postpone revenue generation and create financing and opportunity costs. Source

11. OpenAI, "How agents are transforming work," June 25, 2026. OpenAI reports that agentic AI shifts knowledge work from short interactions toward delegated, long-horizon tasks and that, by May 2026, 70.2% of sampled individual Codex users had made at least one request estimated to exceed one hour of human work. [Source](#)

12. OpenAI, "How OpenAI delivers low-latency voice AI at scale," May 4, 2026, identifying low and stable media round-trip time, low jitter, and low packet loss as concrete requirements for natural real-time voice interaction; and OpenAI, "Build more natural voice experiences with GPT-Live-1 in the API," Sept. 10, 2026, extending full-duplex voice capability into developer applications and business workflows. [Engineering source](#) [API source](#)

13. Anthropic, "Introducing Claude Sonnet 5," June 30, 2026, describing planning, browser and terminal use, and autonomous operation; Anthropic, "Introducing Claude Tag," June 23, 2026, describing access to selected channels, tools, data and codebases and asynchronous work over hours or days; and Anthropic, "DXC will integrate Claude into the systems banks, airlines, and other regulated industries rely on," June 11, 2026. [Sonnet 5 source](#) [Claude Tag source](#) [DXC source](#)

14. Netflix, "Open Connect," describing partnerships with more than 1,000 ISPs to localize substantial amounts of traffic through embedded Open Connect Appliances. Netflix states that each embedded deployment can offload a substantial amount of content traffic from peering or transport circuits and that multiple deployments can be distributed or clustered geographically or across the network to maximize local offload. [Source](#)

15. ETSI, "Multi-access Edge Computing (MEC)," describing cloud-computing capabilities and an IT service environment at the network edge, including fixed access, with ultra-low latency and high bandwidth; and Amazon Web Services, "AWS Wavelength," describing compute and storage in telecommunications-provider data centers for low-latency, resiliency, and edge applications, including ML inference. [ETSI source](#) [AWS source](#)

16. Amazon Web Services, "Use cases for AWS Wavelength," describing healthcare, real-time video analytics, connected vehicles, smart factories, AR/VR, and other low-latency edge workloads; Microsoft, "What is Azure Extended Zones?" (updated June 11, 2026), describing metro- and industry-center deployments for low-latency, throughput-intensive, and data-residency workloads; and Google Cloud, "Google Distributed Cloud," describing data-center and edge deployment for local data processing, survivability, low-latency needs, and AI inference. AWS source Microsoft source [Google source](#)

17. Internet Engineering Task Force, RFC 8655, "Deterministic Networking Architecture" (Oct. 2019). The architecture describes DetNet as dedicating network resources such as link bandwidth and buffer space to flows or classes of flows; identifies resource allocation, service protection, and explicit routes as core techniques; and states that bounded delivery latency is achieved by reserving bandwidth and buffer resources at each DetNet node along the flow path. [Source](#)

-
18. ITU-T Recommendation G.9807.1, "10-Gigabit-capable symmetric passive optical network (XGS-PON)" (Feb. 2023), especially Annex C, Clause C.7, Resource Allocation and Quality of Service. The Recommendation provisions traffic flows with upstream and downstream service parameters; assigns QoS-aware traffic-management functions to OLTs and ONUs based on service specifications and resource availability; and defines fixed, assured, non-assured, and best-effort upstream bandwidth components. It also defines downstream transmission as OLT-to-ONU multiplexing and upstream transmission as OLT-granted transmission opportunities from ONU traffic-bearing entities. [Source](#)
19. Netflix, "Open Connect," describing partnerships with more than 1,000 ISPs to localize substantial amounts of traffic through embedded Open Connect Appliances. Netflix states that each embedded deployment can offload a substantial amount of content traffic from peering or transport circuits and that multiple deployments can be distributed or clustered geographically or across the network to maximize local offload. [Source](#)
20. Amazon Web Services, "Use cases for AWS Wavelength," describing healthcare, real-time video analytics, connected vehicles, smart factories, AR/VR, and other low-latency edge workloads; Microsoft, "What is Azure Extended Zones?" (updated June 11, 2026), describing metro- and industry-center deployments for low-latency, throughput-intensive, and data-residency workloads; and Google Cloud, "Google Distributed Cloud," describing data-center and edge deployment for local data processing, survivability, low-latency needs, and AI inference. [AWS source](#) [Microsoft source](#) [Google source](#)
21. ITU-T Recommendation G.9807.1, "10-Gigabit-capable symmetric passive optical network (XGS-PON)" (Feb. 2023), especially Annex C, Clause C.7, Resource Allocation and Quality of Service. The Recommendation provisions traffic flows with upstream and downstream service parameters; assigns QoS-aware traffic-management functions to OLTs and ONUs based on service specifications and resource availability; and defines fixed, assured, non-assured, and best-effort upstream bandwidth components. It also defines downstream transmission as OLT-to-ONU multiplexing and upstream transmission as OLT-granted transmission opportunities from ONU traffic-bearing entities. [Source](#)
22. Internet Engineering Task Force, RFC 3469, "Framework for Multi-Protocol Label Switching (MPLS)-based Recovery" (Feb. 2003), Sections 3.3 and 3.3.1. The RFC distinguishes equivalent recovery paths from limited recovery paths according to whether adequate resources are available to replace the working path without service degradation. It further describes 1+1 recovery-path resources as fully reserved and 1:1 recovery-path resources as available to preemptible low-priority traffic until a fault requires them for protected traffic. [Source](#)
23. MEF 61, "IP Service Attributes for Subscriber IP Services" (2018), defining Class of Service Names in relation to performance objectives and related Bandwidth Profiles; and MEF 10.4, "Subscriber Ethernet Service Attributes" (Dec. 2018), defining ingress and egress bandwidth-profile flows by Class of Service and bandwidth-profile envelopes that can share and reallocate unused bandwidth among ranked flows. [MEF 61 source](#) [MEF 10.4 source](#)
24. Amazon Web Services, "AWS Wavelength FAQs," describing Wavelength as infrastructure deployed within communications-service-provider networks for workloads requiring low latency, resiliency, data residency, and edge processing; and GSMA, "Operator Platform Telco Edge Requirements - October 22," describing a platform intended to expose network capabilities to enterprise customers and developers so operators can monetize those capabilities, including Quality of Service and low-latency interconnections. [AWS source](#) [GSMA source](#)
25. MEF, "Slicing for Shared 5G Fronthaul and Backhaul" (Apr. 2020). The paper describes deterministic Service Level Specifications that can be strictly enforced for requirements including bandwidth, low latency, and high availability; it defines the Service Level Specification as the technical section of a Service Level Agreement and lists service performance, objectives, metrics, measurement, classification, and bandwidth-profile details among its contents. [Source](#)

26. Big Bang Broadband LLC, "ICQ - Client Executive Overview," Sept. 12, 2026. The source describes ICQ as qualifying physical and optical infrastructure against purpose-specific requirements, identifying capability gaps, preserving unresolved facts and evidence, and supporting decisions involving protection, resilience, diversity, capacity, reserve, future optical capability, capital allocation, and preservation of future choices.

27. CostQuest Associates, "Broadband in America: Fixed Broadband Competition 2025 – January 2026 Edition" (Feb. 3, 2026). The analysis reports that approximately 28% of new fiber construction was reaching homes already served by both fiber and cable, that smaller providers were directing approximately 40% of new builds toward second-fiber markets, and that cable subscriber losses were occurring in newly competitive fiber geographies. [Source](#)

28. MEF, "Slicing for Shared 5G Fronthaul and Backhaul" (Apr. 2020). The paper states that mobile operators without their own mobile-transport facilities use transport services from local service providers; identifies service-provider revenue opportunities in mobile fronthaul and backhaul; and describes mobile-transport service requirements and monitoring involving latency, jitter, bandwidth, and Service Level Specifications. [Source](#)

29. Biarri Networks and Big Bang Broadband LLC, "Type-B and Dual-Parent Protection Analysis Program" SOW (2026), Phase 3 – Test & Cost Comparison. The program compares a best-effort OSP design, that design corrected to Type-B after the fact, and a Type-B-optimized design built from day one, with a three-way BOM and cost comparison intended to isolate the retrofit cost premium. The study is in progress; no result is asserted here.

30. Big Bang Broadband LLC, "ICQ - Client Executive Overview," Sept. 12, 2026. The overview describes ICQ as qualifying physical and optical infrastructure against purpose-specific requirements, preserving unresolved facts rather than replacing them with assumptions, identifying capability gaps, and supporting decisions involving protection, resilience, diversity, capacity, reserve, shared use, future optical capability, capital allocation, and preservation of future choices.

31. Big Bang Broadband LLC, "Infrastructure Capability Qualification (ICQ) Version 6.0," Sept. 2026, Chapter 11, especially §11.13, and Chapter 17, especially §§17.4 and 17.9. The methodology separates candidate generation and technical qualification from later prioritization and capital-efficiency comparison; permits controlled comparison of technically supported alternatives; and identifies evidence-supported consequences including required additions, construction effects, future rework, reserve and optionality effects, and usable cost relationships.

32. Big Bang Broadband LLC, "Infrastructure Capability Qualification (ICQ) - Executive Brief," Sept. 3, 2026. The Executive Brief states the decision sequence explicitly: ICQ establishes the qualified decision population; the authorized client or project authority chooses the applicable business and capital direction; and the authorized engineering-design process converts the accepted direction into engineering design.