BLACKMORE TECHNOLOGY GROUP

BIT TRADE GLOBAL

BTG BRIDGE

A PLAIN-ENGLISH WHITEPAPER FOR INSTITUTIONAL, INDUSTRY AND COMMUNITY REVIEW

BUILDING A KNOWLEDGE CAPITAL ECONOMY

How an interconnected system for research, executable invention, lasting knowledge, lawful rights, commerce and investment is being engineered in British Columbia

October 2026 | Grand Forks, British Columbia, Canada

Verification. Usability. Legitimate rights. Productive knowledge.

CONTENTS

Executive overview: the whole idea in one picture
Part I. Understanding the problem BTG is addressing
Part II. The technologies: what each system is and how it actually works
Part III. The seven transformations as a connected economic system
Part IV. One complete example: following an invention through every system
Part V. What the completed architecture means for industry and everyday life
Part VI. The economic opportunity and how BTG could earn revenue
Part VII. Grand Forks, British Columbia and the wider public interest
Part VIII. Trust, regulation and practical boundaries
Part IX. The demonstrated engineering foundation and the complete system being built
Part X. What this infrastructure could become
Part XI. Questions a reader may reasonably ask
Conclusion. A different relationship between knowledge and economic opportunity
Appendix A. Plain-English guide to the named systems
Appendix B. Selected sources and evidence

Blackmore Technology Group Limited
Grand Forks, British Columbia, Canada
October 2026
Whitepaper for institutional, community, research, industry and public-interest review

A different starting point for economic growth

For much of the modern economy, a company comes first. It raises money, hires people, gathers information, develops an invention or product, and eventually earns revenue. The value of the information and inventions it creates is normally recognized through the value of the company that controls them.

Blackmore Technology Group Limited (BTG) is engineering an additional path: a particular piece of useful knowledge - a lawful dataset, a tested algorithm, a reusable software capability, a research result or a reconstructable invention - can have its own lasting identity, demonstrated usefulness, defined legal permissions and commercial history. A business may then be built or financed around that asset. The asset does not replace the company or the law. It becomes a more visible starting point for enterprise, collaboration and investment.

The name BTG has two distinct uses in the work described here. **Blackmore Technology Group Limited** is the originating company and technology developer. **Bit Trade Global** is the name of the asset-rights and proposed investment-market infrastructure within the wider system.

This paper explains the technologies from the ground up, using ordinary language and practical examples. It describes the integrated engineering architecture, the economic opportunity, the opportunities for Grand Forks and British Columbia, and the safeguards that matter when knowledge begins to move between autonomous machines, organizations and markets. It also identifies concrete research demonstrations without portraying an evolving system as a publicly operating securities exchange.

Executive overview: the whole idea in one picture

Imagine a manufacturer in the Boundary region has equipment that occasionally stops in a way its ordinary software cannot predict. A research system studies permitted machine data, devises a new detection method, writes working code, tests it, discovers a weakness and repairs it. The useful part is kept in a form that another business can run without buying an entire application or surrendering its own data. The history of the experiments travels with the capability. A rights system states who can use it and under what conditions. A customer pays to use it. If researchers or data contributors are entitled to payment, the money can be accounted for and distributed. A legally organized business may then seek capital to improve the invention. Revenue or properly authorized financing may pay for the next round of research.

BTG is engineering the systems needed to connect that entire sequence. The combination is the Knowledge Capital Economy.

To understand it, think of a modern industrial city rather than a single computer program. Research laboratories discover new methods. Libraries preserve what is learned. Workshops manufacture reusable components. Assembly plants connect those components into complete machines. Property offices record claims, contracts and permissions. Inspectors check whether things work and whether the evidence can be trusted. Commercial marketplaces connect products with customers. Banks transfer actual money. Financial markets can help qualifying businesses raise capital. A shared record connects events without letting one office pretend to do the job of another.

In the BTG architecture, those roles are carried out by named systems:

- **NIKI CGSI-SI:** the invention and research intelligence, designed to investigate unfamiliar problems, create and repair executable methods, and learn from verified results.

- **Computational FarmNodes:** separate digital laboratories in which research and testing can take place, and whose verified discoveries can be shared between authorized participants.
- **BTDU:** the knowledge system that preserves authoritative information, its meaning, its relationships, its origins and the evidence of what worked or failed.
- **DataChips and AI Chips:** small, identifiable, usable capabilities - calculations, algorithms, data operations or model-based functions - with known inputs, outputs, versions and permissions.
- **Circuits:** working assemblies made by connecting multiple DataChips or other permitted capabilities.
- **Circuit Genomes and Farm Genomes:** versioned plans that describe how a circuit or even a whole research laboratory is built, recreated and improved.
- **DataChip Composer and the Universal SDK:** the developer environment and integration tools used to turn ordinary code into compatible components and combine components into applications and services.
- **ENTITY:** the system for lasting asset identities, provenance, ownership claims, licences, permissions, contribution relationships and economic entitlements.
- **ADAM:** the governed decision and state-change engine that determines whether a proposed action is authorized and commits an accountable result.
- **Universal Transaction Fabric:** the record and messaging fabric that connects signed events across independently operating systems.
- **Verifiable Reality:** the evidence framework that keeps a claim, a laboratory result, a signed contract, an independent assessment and an actual payment from being mistaken for the same thing.
- **Bit Trade Global:** commercial asset discovery and rights exchange, connected to a distinct, appropriately authorized investment infrastructure for financial interests in asset-related businesses.

The systems are not interchangeable. NIKI can invent a method but cannot create someone else's legal ownership. BTDU can retain the source but does not make it profitable. ENTITY can record what a contract permits but cannot manufacture a missing contract. ADAM can enforce an authorized transition but cannot replace a bank or court. A wallet can display an entitlement but cannot, by itself, create a valid share in a company. **The design becomes trustworthy by connecting those functions while preserving the real authority of each.**

The seven transformations are the economy-level changes these systems jointly make possible: continuous invention; software as reusable intellectual assets; portable rights; lasting knowledge capital; financing around useful assets; fair contributor participation; and a cycle in which productive knowledge helps finance more productive knowledge.

Part I. Understanding the problem BTG is addressing

1. Useful knowledge is difficult to recognize as a separate asset

A business can own a valuable dataset and still struggle to describe exactly why it matters. A small research team can create a sophisticated algorithm but lack the marketing capacity to sell a full software product. A municipality can have valuable operational observations while being obligated to protect sensitive information and public ownership. A software company can rely on dozens of useful internal functions that nobody outside the company can purchase separately. A scientist can publish a successful experiment without preserving the exact working software and failed experiments that made the result credible.

The world has many solutions to pieces of this problem: code repositories, databases, scientific archives, patent offices, software libraries, cloud services, data marketplaces and stock exchanges. BTG is not claiming that it invented these institutions. Its architecture addresses the missing continuity between them.

The recurring questions are simple: What was created? Can another party reproduce it? Which evidence supports its usefulness? Who may lawfully use it? What other work went into it? How does it interact with

existing software? Can a customer buy the benefit without taking all the underlying data? If money is earned, who is entitled to it? Can a business seek investment against the genuine productive potential of the asset?

Ordinary software systems often answer these questions in different, disconnected databases. BTG's approach is to give the intellectual asset a recognizable identity throughout its life, from research experiment to software component, authorized use, commercial transaction, contributor statement and, where applicable, legally structured investment.

2. What 'data as an asset' means - and what it does not mean

Data can be a record of a machine temperature, a soil test, a mapping observation, a research measurement, an equipment failure or a mathematical result. On its own, a copied file may be easy to reproduce and have little commercial value. The economically important questions are whether the information is lawful to use, accurate, difficult or expensive to reproduce, relevant to an important problem, timely, and capable of improving someone's work.

Not every item that contains data is something its holder can sell. Information can be covered by privacy rights, contractual restrictions, Indigenous data sovereignty, public-sector duties, confidentiality obligations and intellectual-property law. BTG's use of the word *public* refers to making eligible asset descriptions, rights, evidence and permitted market information discoverable. **It does not mean publishing raw personal files, hospital records, private locations or confidential research.**

The architecture also distinguishes three economic transactions. A business may pay to **use** a dataset or run an algorithm. A person or organization may receive **contractual proceeds** because it contributed something under a valid agreement. An investor may buy a **lawfully issued financial security** in a company or other qualifying structure associated with an asset. Buying access to an algorithm is not the same as buying shares, and receiving a royalty is not automatically becoming an investor.

3. What makes BTG's approach different from a simple AI or data marketplace

An AI marketplace ordinarily provides models or AI-powered services. A data marketplace ordinarily lists information or access products. A software component repository distributes code. An investment platform connects issuers with money. BTG's architecture is broader because it joins all of these roles to a persistent record of knowledge, executable functionality, evidence and authority.

A successful invention need not stay inside the laboratory that produced it. It can be reconstructed, packaged as a controlled capability, checked against its source and rights, assembled with other capabilities, used under licence, assessed for commercial usefulness and - through a separate lawful financial structure - become part of a business's capital-raising case. A failed experiment also has a place, because a known failure can help another researcher avoid wasting resources.

The distinguishing feature is not simply that data may be bought or sold. **It is the continuity between how useful knowledge is produced, what proves it works, which rights survive its reuse, and how economic value can support further invention.**

Part II. The technologies: what each system is and how it actually works

4. NIKI CGSI-SI: the invention engine

What NIKI is

NIKI is BTG's cognitive research and invention system. In ordinary terms, it is being engineered as a research worker that can do more than answer questions or produce a paragraph of computer code when prompted. Its job is to accept an authorized problem, find relevant permitted knowledge, devise candidate solutions, express them as working computations, run bounded experiments, identify mistakes, improve its designs and retain what the evidence justifies.

A conventional chatbot may give a plausible suggestion. NIKI's architecture is concerned with a more demanding question: **Can the system create a working method that was not simply handed to it, test that method against an independent standard and preserve the exact result for other work?**

How NIKI works within BTG

A human or authorized research organization begins with a **research mandate**: the problem to study, data it may use, boundaries on its actions, computing budget, technical acceptance test and authorized decision-makers. For a manufacturing problem, the mandate might be to reduce false alarms when several sensors disagree. For forestry, it might be to improve prediction of a particular equipment failure. For agriculture, it might be to detect a pattern in permitted moisture and weather readings.

NIKI consults authorized knowledge from BTDU and other specified inputs. It can generate a candidate algorithm, a working program, or a proposed arrangement of DataChips. The exact generated bytes and version of the source are retained. The proposed program is treated as untrusted: it runs inside a restricted environment with limits on memory, processor time, processes, files and outside connections. An independent test compares outputs with a known answer or other measurable standard. A failed test is recorded, not rewritten into a success. A repair is recorded as a new version with a visible relationship to the failed one.

If a candidate passes the relevant tests, the evidence travels onward. Its source, compiled executable, input and output definitions, dependency versions, resource limits, observed reliability and permitted-use status can become part of a governed DataChip or Circuit. NIKI may help investigate the next problem, but it does not approve its own legal rights, independent test findings or financial transactions.

Why it matters

The economic importance lies in the possibility of producing more **useful, verifiable intellectual work** for a given amount of time and resources. Research organizations normally employ specialized people to explore methods, write code, test, correct and document results. NIKI is designed to make portions of that process repeatable and available to many participating laboratories, while retaining responsibility and independent technical review for consequential results.

NIKI does not need to become an all-knowing intelligence to be valuable. If it consistently produces qualified improvements to specific industrial, scientific or software problems, each useful result can become a separate asset with a known history. Large numbers of those results could make the wider economy more productive. Whether that advantage exists at commercial scale is established through independent comparisons of useful results, costs, failures, human involvement and customer outcomes, not by calling every piece of generated code an invention.

Example anyone can follow

Suppose four sensors on a piece of machinery report possible faults. One sensor may be missing, another may be invalid, and the others may agree. A proposed algorithm must interpret every combination correctly. NIKI first produced a version that incorrectly treated certain missing readings as agreement. The error was exposed by a fixed test. The repaired source, which retained the original chain of evidence, later produced the correct outputs across all 65,536 possible inputs in recorded Windows and Linux WebAssembly execution. This is an example of a real local research and repair chain, not a claim of general-purpose scientific superiority or a completed commercial product.

5. Computational FarmNodes: laboratories that exist as software systems

What a FarmNode is

In this architecture, a **FarmNode is a whole logical research laboratory**. It is not simply a computer connected to a network and not merely a GitHub repository. A laboratory can contain a research agenda, permitted datasets, tools, programs, NIKI workers, DataChips, Circuits, experimental procedures, computing limits, evidence records and rules describing what its operators may share.

A FarmNode may run on one machine, across cloud infrastructure or through a cluster managed by an organization. The laboratory has its own identity even if it moves to different physical hardware. Its repository may contain versioned program code, but large datasets and active computing services belong in appropriately controlled storage and runtime systems.

How farms collaborate

A research farm can receive an authorized problem, consult BTDU knowledge, run NIKI experiments and package the resulting evidence. Another farm, controlled by a different research organization, can receive a permitted research capsule: exact source references, input descriptions, proof of origin, rights restrictions, evaluation instructions and what may be disclosed. It can independently recreate the test or build on the qualified result. A third farm might investigate an edge case or develop a compatible capability for a different industry.

Each farm records **what it observed itself**. An encrypted delivery acknowledgment means information arrived; it is not a claim that the recipient independently recreated the invention. Three processes on one desktop can demonstrate message routing but are not three independent institutions. Genuine cross-institution verification depends on separately controlled identities, hosts, evaluators and evidence.

Farms as entire virtual research worlds

A farm can represent a robotics laboratory, a materials research facility, a simulated agricultural environment, a manufacturing test centre or a specialized mathematics laboratory. It may contain several interacting Circuits and an evolving **Farm Genome** describing the equipment, procedures, data rights and resources needed to recreate the laboratory's working arrangement. That is why the farm concept is broader than the ordinary definition of a distributed-computing node.

Thousands of such laboratories could specialize. One farm works on machinery. Another focuses on weather models. A third tests mathematics. A fourth verifies software quality. Where rights and technical formats permit, qualified findings can be reused across these environments. The same small invention could be useful in several industries, and new experiments could generate additional capabilities without each organization starting from nothing.

Why farms matter to a rural region

A physical laboratory may require expensive property and equipment. Some computational research, by contrast, can be performed remotely, although it still requires qualified staff, computing resources, secure infrastructure and

professional oversight. That opens a path for smaller centres such as Grand Forks to participate in technically sophisticated work through local teams collaborating with institutions elsewhere. A university or manufacturer elsewhere can assess the work without assuming that a rural address makes it less credible; what matters is the reproducible evidence and capability.

6. BTDU: the knowledge system and durable memory

What BTDU is

BTDU is BTG's knowledge representation, storage and reconstruction foundation. It is not just a directory of files or a chatbot's conversation history. Its purpose is to preserve a reliable original record while allowing many useful relationships, interpretations and working views to be constructed from that record.

BTG summarizes this principle as **'one authoritative payload, unlimited semantic views.'** The payload is the content that is treated as the authoritative source for a particular record. A semantic view is one way to understand or use that content: a set of words and meanings, a research relationship, a search index, a machine-readable calculation, a summary or a connection to an experiment. Many views can refer to the same source without each becoming a competing master copy.

Atoms, bonds and the Atom Universe

The BTDU architecture includes an **Atom Universe**. In everyday terms, atoms are fundamental identifiable pieces of information - for example symbols, numbers, letters and other precisely represented units. Bonds describe how pieces are joined or related: the characters in a word, the structure of a formula, the dependencies in a piece of code, the relationship between a scientific claim and its supporting measurements, or the ancestry of a derived asset.

This is different from saying that ordinary databases cannot store relationships. The BTDU design seeks a shared, reconstructable representation that can connect language, software, mathematics and research evidence while preserving where the information came from. Meaning still depends on context. A word appearing in a scientific report is not proof that the report's conclusion is true, and a connection in a knowledge graph is not automatically a legally owned asset.

Reconstruction and evidence

Suppose the exact source code of an algorithm is preserved in BTDU. A scientist later wants to recreate what was tested. The system must retrieve the correct version, its inputs, its expected outputs, the source of the measurements and the limits of the test. The derived search and semantic views help locate and interpret the material. The authoritative payload is used when exact bytes matter. If a summary says something different from the experiment record, the record prevails.

Failed experiments are important. A study that establishes why a method does not work may prevent repeated expense. Contradictory findings must remain visible instead of being silently erased. Rights also travel with the data: knowing that a source exists does not automatically give a user permission to download it, train a model on it or resell it.

BTDU and AI

AI systems often struggle with stale information, invented citations and forgotten context. BTDU provides a structure in which NIKI can reference pinned source versions, experimentally observed results and permitted knowledge relationships. That improves the ability to audit an AI result; it does not magically guarantee the AI reasons correctly. Reconstruction claims, compression claims and performance claims must each be tested against real baselines.

Why accumulated knowledge becomes capital

A new discovery can save resources immediately. Its retained methods, input data rights, failures and relationships can make later discoveries less expensive. A knowledge system adds economic value when it measurably reduces duplicate research, improves decisions, speeds correct reconstruction or creates new legitimate commercial uses. The value is **the productive work the knowledge enables**, not the number of words, atoms or bonds stored on a disk.

7. DataChips and AI Chips: turning capabilities into identifiable building blocks

The simplest explanation

A **DataChip is a reusable piece of working software with a clearly defined purpose**. Its name may suggest a tiny physical semiconductor, but it is a digital capability. It can calculate a result, transform permitted data, evaluate a condition, execute a scientific method or carry out part of a larger process.

An ordinary software application can contain hundreds of useful routines. Consider a navigation application that estimates distance, determines wind exposure, checks terrain constraints, compares possible routes and displays a result. Normally the customer buys the whole application or pays to access an entire service. With DataChips, each meaningful capability can be identified, tested, updated and reused in other software where the rights allow.

What a DataChip carries

A DataChip is not merely a function name or a box on a screen. Its description includes the particular task it performs; the information it requires; the result it returns; the version of the algorithm; what operating environment can run it; what computing and security resources it requires; what independent tests have established; which other software components it depends on; and what legal permissions apply to its use.

For example, a DataChip might accept a machine's last five temperature measurements and return a fault score. Its public description can explain that function. Its test evidence can indicate which types of machinery were studied. Its rights passport can limit use to manufacturing maintenance, prohibit resale of the source code and set the expiry of a licence. A buyer might be permitted to call the computation without being allowed to see private training data.

The **DataChip Composer** is designed to create and connect these components using both ordinary programming code and a visual arrangement of functional blocks. A developer is not required to abandon established programming languages to participate. Existing code can be imported or wrapped as a candidate DataChip, then qualified against its real behaviour and source licences. Visual composition helps people understand and assemble compatible functions; it does not substitute for testing or software engineering.

What an AI Chip adds

An **AI Chip** uses a trained model or another AI capability. Unlike a simple deterministic calculation, some AI outputs can vary or be uncertain. The component therefore identifies the model, its version, applicable training and inference rights, supported inputs, expected performance, limitations and human review conditions. A company can use a model-backed component and a deterministic component within one larger system, provided each has the right safety controls.

Why standalone capability matters economically

A small engineering firm may invent one excellent calculation without building a complete smartphone app or cloud platform. A large manufacturer may need only that calculation and want to embed it into existing equipment. A software developer may improve it for another industry. When a capability has a stable identity, reproducible behaviour and defined rights, these parties have something precise to evaluate and contract for. It can become a commercial asset rather than an invisible feature buried inside a larger product.

8. Circuits: building complete working processes from DataChips

A **Circuit is a connected group of working capabilities**. A simple way to picture it is an assembly line. One station receives a measurement. The next cleans or checks it. A third makes a decision. A fourth records the result. Each station has a specific job, but the complete line accomplishes more than any station alone.

In a digital Circuit, the connections are not physical wires. They are explicitly declared routes between inputs and outputs. A DataChip's output must be suitable for the next component's input. When the units or expected formats differ, the system must use an approved conversion rather than silently feeding incompatible data to the next step. Circuits can include ordinary algorithms, AI Chips, access controls, logging operations and permitted remote services.

Consider a wildfire-planning research Circuit using publicly available or otherwise authorized inputs. The first DataChip interprets weather measurements. A second calculates a terrain condition. A third compares route constraints. A fourth labels uncertainty. A final reporting component assembles the result for a qualified human planner. The complete Circuit is more powerful than any one function, but the decision remains a planning aid rather than an autonomous command to emergency workers.

Why the original source of each part still matters

A Circuit should not hide what it depends on. The use of one component may be allowed for internal research but forbidden for public redistribution. Another may require additional testing in a safety-critical setting. A third may be owned by an unrelated university. Joining them does not erase their legal or safety requirements. The assembled Circuit carries the relevant lineage and effective restrictions of its parts.

The execution engine

The **CircuitRuntime**, part of BTG's Recursive Fabric work, is the software that actually executes the connected operations. It knows which operation comes first, how outputs are passed along, which errors stop the process, how resource limits are enforced by supporting systems and what evidence is recorded. A Circuit specification is the wiring plan; the runtime is the machine doing the work.

That difference is visible in BTG's engineering evidence. Early work ran a digest-processing Circuit connected to NIKI's research records. Later local engineering connected the actual NIKI-generated WebAssembly sensor algorithm as a typed operation within a two-node Circuit, followed by a SHA-256 digest operation. The recorded output matched the fixed oracle across all 65,536 inputs. These are different stages of engineering, and the later result is more directly relevant to executable composition.

9. Circuit Genomes: preserving the design of a complete capability

A **Circuit Genome is a versioned recipe for building and rebuilding a Circuit**. A recipe is more than the finished meal: it explains the ingredients, their quantities, the steps, and what must change when one ingredient is different. Likewise, a Circuit Genome captures the identities and versions of its component DataChips, their connections, the accepted input and output definitions, dependencies, relevant tests, rights and reconstruction instructions.

Why call it a genome? In biology, a genome provides an organized set of instructions involved in forming and maintaining a living organism. BTG uses the analogy for a working digital architecture that can be recreated and revised. It is **not a living organism**, nor does the name mean the software spontaneously becomes intelligent. It describes a structured design whose versions can be inspected and whose changes can be measured.

Suppose a manufacturer uses a Circuit to identify developing bearing faults. A new measurement method becomes available. An updated Genome may replace one DataChip, add a verification step, or change how multiple sensors

are combined. The previous qualified version remains recoverable. The new arrangement must be tested as a new version rather than inheriting every claim about the previous one automatically.

This opens two different forms of commercial value. One company might licence a particular DataChip - a single useful calculation. Another might licence an entire Circuit Genome - a tested configuration that combines several capabilities to solve a more significant problem. NIKI and FarmNodes can investigate new combinations, but any new design must earn its own reproducibility and rights evidence.

The related Farm Genome

A **Farm Genome** works at a higher level. Instead of describing one functional process, it describes how to assemble a complete virtual research laboratory: the kinds of research it performs, instruments and tools available, datasets it may access, its experiment procedures, permitted capabilities, evaluation methods and operating constraints. That makes a FarmNode reproducible as a research environment, not merely a single program.

These two genomes should never be confused. A Circuit Genome builds a particular executable solution. A Farm Genome builds a laboratory capable of producing and evaluating many solutions.

10. ENTITY: the identity, rights and economic relationship system

What ENTITY is

ENTITY gives a durable identity to people, organizations, software capabilities, datasets, inventions and related economic rights. It records the relationships that matter when knowledge moves between systems. A DataChip may travel from one cloud service to another while remaining the same underlying version. A researcher may grant a manufacturer permission to run it for a limited purpose. A contributor may be entitled to part of a particular commercial revenue stream. ENTITY is designed to preserve and verify these distinctions.

Imagine a land registry combined with a library catalogue and a permissions desk. The catalogue identifies an item. The permissions desk determines who is allowed to use it and in what way. The registry preserves claims and supporting documents. But the office does not become the law itself: if a legal title is disputed, contracts, legislation and relevant authorities must decide the outcome.

Provenance: knowing where something came from

Provenance means the origin and history of an item. In BTG, an asset record can point to its original research project, source data, code version, parent DataChips, test evidence and contributing organizations. Provenance answers, 'Where did this come from?' rather than merely, 'What is this file called?'

A cryptographic fingerprint can establish that two versions of a file contain the same bytes. A digital signature can establish which key signed a statement. Neither is automatic proof that the signer owns the copyright, has the right to sell the dataset or can issue shares. ENTITY therefore distinguishes a technical claim from documentary legal authority and recognizes conflicts, expiring permissions and revocations.

Rights Passports

A **Rights Passport** is a readable and machine-verifiable description of what permissions are associated with an asset. It might say that a buyer may execute a particular algorithm for one year, only inside its manufacturing systems, with no permission to download underlying private data, modify the original source or sublicense it. Another agreement might allow research use but forbid commercial use. An assignment of intellectual property, a commercial licence and a share in a company have different legal meanings; a generic 'owner' label cannot replace those meanings.

Before a protected action takes place, a participating service checks the relevant rights, identity, time period, purpose, version and restrictions. If an entitlement expires or is lawfully revoked, a new unauthorized action

should be blocked. The historical audit remains available as law permits so parties can determine what happened when a contract was still active.

Contributor relationships and sovereignty

ENTITY also relates assets to their creators and contributing sources. The system can represent multiple parent assets, companies, researchers, data contributors and agreed allocations of commercial proceeds. It can identify that a derived invention depended on a particular licensed dataset without automatically asserting that every source contributor owns the derived invention. Contributions and economic entitlements are established under actual agreements.

Organizations and Indigenous Nations may have distinct data-governance requirements, and some information is not appropriate for commercialization at all. The system is designed to enforce purpose and jurisdiction boundaries rather than treating everything that can be copied as freely transferable.

Why ENTITY matters beyond a single BTG service

An important distinction is portability. If the asset, its right and its evidence are legible only inside one company's private database, participants depend entirely on that company. ENTITY's role within the wider architecture is to support interoperable verification so legitimate rights can remain meaningful as assets are reused by different applications, farms and commercial partners. Its design is not to impose an automatic fee on every unrelated party merely because an asset has an ENTITY record. Service fees and ownership income must come from the appropriate contracts.

11. ADAM: the controlled engine for decisions that change the system

ENTITY describes identities and claims. **ADAM determines whether a requested action is permitted and commits an accountable change to the relevant system state.** A state is simply the system's current agreed record: a licence is active, an asset version is qualified, a research budget is approved or a transaction request is on hold.

To understand the difference, imagine a business keeps a secure employee-access register. ENTITY identifies an employee and their job permissions. ADAM is the gate that checks a particular request and either permits a recorded change or refuses it. The fact that the employee once had permission does not mean that every later request is automatically allowed.

A NIKI agent might request that a completed experiment be retained. ADAM checks the appropriate authorization, the exact asset version, the owner's current rules and whether the record has changed since the request was prepared. If valid, the authorized domain makes the change and records its receipt. If stale, disputed or unauthorized, it rejects or holds the command. An attempted replay of a completed transaction must not create a second payment or second issuance merely because a network message was sent twice.

ADAM is central to autonomous systems participating in real business activity. Intelligence can recommend an action; a governed state engine decides whether that recommendation can produce a real side effect. That is safer than handing an AI agent unrestricted access to a company's databases, funds and legal records.

ADAM is not the registrar of securities, a bank, a judge or a substitute for the authorization of a human rights holder. It enforces rules backed by legitimate authority within its scope.

12. Universal Transaction Fabric: connecting the record of what happened

The **Universal Transaction Fabric** is the communication and event-recording layer connecting activities across the BTG systems. In everyday terms, think of it as the system of numbered receipts and dependable delivery between different departments. Each event can identify what happened, which asset or right it concerned, which authority acted, which rules were checked and how to retrieve the relevant evidence.

When NIKI produces an experiment result, the research service records it. When a DataChip version is qualified, the capability service records that event. When an owner grants a licence, the rights system records an authorization. When a customer actually pays, the financial service references the external payment evidence. A public transparency service may disclose a safe verification summary without publishing a private customer's identity or contract.

This fabric is designed to tolerate the ordinary failures of distributed computing. A message may be delayed, delivered twice or arrive after an outage. The architecture uses stable references, signatures, version checks and duplicate handling so repeated messages do not change legal balances twice. When one service fails after another has acted, the recovery process must reconcile the authoritative records rather than simply declaring all parts of the transaction successful.

The name 'transaction' has a broad meaning here: it may refer to a confirmed technical event, a licence permission or an actual financial transaction. **A technical event receipt is not money.** It is also not automatically a corporate share certificate or enforceable assignment of intellectual property. Each real-world authority remains separately identified.

13. Verifiable Reality: telling proof from belief

A major problem in AI and digital markets is that different kinds of claims can look equally convincing on a screen. A program says an algorithm works. A test log shows it worked in one environment. Another laboratory independently reproduces the result. A rights holder signs a valid licence. A bank confirms a payment. A regulator issues a written market authorization. These statements have very different meanings.

Verifiable Reality is the evidence-classification approach that keeps those meanings separate. It is designed to give a reader or another system the evidence behind an assertion and the proper degree of confidence in that assertion. A cryptographic hash can verify the integrity of source code; it does not prove the algorithm is useful in an actual factory. An internal pass report can document a laboratory's work; it does not certify an independent institution performed the same test. A signed contract can confer a defined permission; it does not make a disputed claim to unrelated property valid.

For public and institutional adoption, this matters enormously. A manufacturer wants to know which test supports a promised efficiency gain. A public body wants to know whether private data is legally available. An investor wants to distinguish independent valuation from a seller's estimate. A contributor wants to know whether money was actually received before a royalty is shown as payable. Verifiable Reality helps those questions remain answerable without treating AI-generated explanations as evidence by themselves.

14. DataChip Composer, the Universal SDK and BTG Command: the systems people use to build and operate it

Most of the named technologies describe engines and registries. Developers and operators also need actual working tools.

DataChip Composer is BTG's development environment. It brings conventional source-code editing and graphical composition together. A developer can create or import a module, define its inputs and outputs, connect it to existing DataChips, inspect dependencies, assemble Circuits and prepare versioned packages. The visual layout helps reveal how a complex system is constructed, while ordinary code remains available for capabilities that cannot be adequately expressed with drag-and-drop blocks alone.

The Universal SDK provides integration points so different applications and devices can use compatible BTG capabilities. A smartphone app, a Windows workstation, a cloud service or an industrial machine may have different operating systems, hardware resources and security constraints. The SDK's role is to make the supported contracts understandable across these environments, not to pretend one unmodified executable can run on every

device. Open interfaces, including appropriate WebAssembly components and typed interface definitions, help reduce dependence on a single vendor's environment.

BTG Command Centre is the authorized engineering, device-diagnostics and support environment. It is designed to connect BTG applications and permitted devices to real logs, build information, tests and controlled support sessions. That function is important because reliable products require troubleshooting and evidence after deployment. A support console should not have automatic unlimited authority over someone else's data, legal rights or bank accounts. Access remains authenticated, scoped and auditable.

These tools matter commercially because a powerful architecture nobody can operate is not a usable platform. The point is to let developers assemble actual products, let businesses integrate the capabilities they purchase and let authorized engineers diagnose problems without erasing the boundaries between software support, confidential data and financial authority.

15. Bit Trade Global: the commercial and investment economy around knowledge

One name, two legally different markets

Bit Trade Global is the financial and commercial arm of the overall Knowledge Capital Economy. It is built around an idea that sounds simple but has major implications: **the useful underlying knowledge asset should be visible as an economic subject in its own right**. A company may own or control valuable rights in an invention; investors normally value the company as a whole. BTG is engineering a way to describe the particular asset, its provenance, technical usefulness, customers, restrictions, commercial income and possible contribution to an enterprise's value.

The first market is for **commercial rights**. A buyer might pay to run an algorithm, access a permitted dataset, receive a restricted analysis, use a qualified Circuit, or licence particular software. These transactions are governed by intellectual-property, contract, data-protection, tax and other applicable law. They are not automatically financial securities.

The second is for **investment instruments**. Here the buyer is an investor rather than a user of the software. An eligible legal issuer may, with the proper corporate and securities arrangements, offer shares or other defined financial interests in a business associated with an asset. Investors receive exactly the rights stated in those instruments, such as an interest in company equity or an authorized income arrangement. They do not automatically receive ownership of a dataset or permission to read confidential data.

The distinction is essential. A person who pays to run a robotics algorithm has bought permitted utility. A person who owns shares in the company commercializing that algorithm owns the corresponding legal financial interest. A contributor with a contractual share of licensing revenue holds a different entitlement again. The system records all three without inventing a single token that means everything.

How an asset becomes discoverable

Consider an algorithm used to reduce waste in industrial production. Before a commercial listing, its records can identify its version, origin, qualified use case, important dependencies, permitted rights, supported customer environments and available technical evidence. A buyer can see a non-sensitive summary without gaining access to protected source material or private training data.

The marketplace can support **a Rights Passport** that makes the offer understandable: what may be used; by whom; in which industry; for how long; with what quantity or operating limits; whether the buyer can modify or sublicense it; and what happens if the right expires. For a marketplace serving many organizations, such clarity is often more important than a sophisticated trading chart.

A buyer's transaction follows the agreement. A valid customer and licence are identified; the price and terms are accepted; authorized money channels settle any payment; the defined capability or access is delivered; the

transaction and relevant receipts are retained; and any applicable contributor allocation can be calculated. Where licensed access is remote rather than a file download, the customer may receive a result without ever receiving the underlying private source.

What an asset valuation means

A qualified invention may have several different economic values. It can cost a certain amount to reproduce. It may be expected to earn future licence income. Buyers may make genuine offers to use it. A comparable invention may have sold under different terms. An algorithm may produce independently measured savings for a customer. If a business issues a security and that security trades lawfully, there can also be an observed market price for the security.

These are related facts, not interchangeable prices. **An asset's estimated value is not automatically its company's share price, an investor's return or the accounting value that may appear on a balance sheet.** Nor does a high estimate create actual cash for the inventor. A mature marketplace needs to preserve the method, assumptions, evidence and date of each valuation. The market's credibility depends on separating an independent estimate from a seller's marketing statement or an AI-generated projection.

Why asset-first financing changes company formation

The central economic inversion begins with a qualified knowledge asset rather than requiring a large company before the asset can be recognized. Imagine a new robotics recovery method has a traceable invention record, documented legal rights, independent technical evidence and customers willing to buy licences. Those facts may allow a new operating business to raise capital to improve, service and commercialize it.

An issuer might hold several assets, or the economic business case might focus on a specific invention. In either case, investors need disclosures about the actual company, its debts, managers, dilution, expenses, security rights, customers, risks and control of the underlying technology. Investors do not get automatic rights to copy the underlying data simply because it contributes to the value of the issuer.

The system is being engineered with Canadian securities requirements in mind. A lawful restricted primary offering, a secondary trading system and a recognized securities exchange are different activities with different legal obligations. **Engineering an order book is not the same as obtaining permission to operate a stock exchange.** Any live securities-market functionality requires the relevant legal, operator, dealer, custody and regulatory arrangements. The infrastructure can be designed in advance without representing itself as already authorized.

DCOs, tickers and the wallet

BTG uses **DCO** to identify and structure individual digital commercial or intellectual assets and their defined associated rights. The working portfolio includes a proposed robotics failure-and-recovery dataset (**DCO-000001**) and a separate BTG-controlled robotics algorithm (**DCO-000002**). The distinction matters: owning a dataset is not necessarily the same as owning all algorithms developed from it; rights in the source data and derived software must each be established.

An asset identifier or ticker makes something findable. It is a name or record handle, not a magical source of legal scarcity or title. Two businesses can sometimes licence the same information non-exclusively. A company share, by contrast, is a particular legal unit of an issuer with specified rights. Those concepts must remain separate.

The **BTG wallet** is designed to show the user's eligible assets, commercial entitlements, contracts, contributor interests and - where properly issued - financial holdings. It can help display a Rights Passport, an accepted trade receipt, a balance or an asset transfer status. It remains a view onto the authoritative legal and financial records. The legal share register and bank or settlement provider determine facts that a wallet interface alone cannot establish.

The role of money and transaction proofs

The wider economy uses fiat currency and approved settlement rails rather than requiring every participant to buy a speculative BTG coin. A signed transaction-proof token can attest that a permitted contract or event occurred. It may support public transparency through privacy-preserving commitments, including a public verification page, while leaving confidential contracts, personal identities and banking records protected. The token should never be presented as money simply because it carries a digital signature.

16. Contributor participation: how people and organizations can share in legitimate value

The Knowledge Capital Economy is not just about paying technology companies for access to information. It also addresses people and organizations whose authorized contributions help create productive assets.

The word *contributor* can mean different things. A research institution may supply a lawful dataset. An engineer may contribute source code under an employment or licensing contract. A group of inventors may develop a method together. An individual may voluntarily permit a narrowly described use of personal information. A cooperative or Indigenous institution may manage a collectively governed source subject to its own legitimate authority. These cases cannot be treated as though everyone holds identical ownership rights.

From permission to a documented economic entitlement

A sound contribution arrangement begins with the precise right that can be granted. The agreement identifies the source, permitted purpose, duration, parties, privacy restrictions, ability to withdraw future permission where applicable and the economic return, if any. An entry in a database is not a substitute for consent or legal title.

The system can connect an agreed contribution to the assets that use it. If a manufacturer pays for a licensed application, and the agreement provides a portion of qualifying income to named participants, the payment can be traced to the specific receipt and allocation rule. The system must account for disclosed fees, taxes, refunds, other royalties and reserves before distributing the amount actually available under the contract. The combined claims against one receipt must not exceed the money the agreement makes distributable.

A straightforward example

Suppose a research collective and a machinery manufacturer jointly create a rights-cleared fault database. Their agreement provides that after specified expenses, the remaining eligible licensing proceeds are divided among the collective, the manufacturer and named contributors. A customer pays an invoice for a permitted use. The payment provider confirms settlement. The financial system calculates the shares using the actual receipt and agreed rules. Each participant receives a statement showing the basis for the allocation, pending payments and any withholding. If a refund later occurs, the adjustment is recorded rather than silently disappearing.

If a later Circuit uses the same dataset through three different component paths, that should not automatically triple the contributor royalty. The economic claim follows the legitimate contract and revenue allocation, not the number of times the asset's name appears in a software graph.

Privacy and fairness

Some data should not be commercialized, no matter how technically attractive it appears. Information about children, patients, precise personal movements, private communications or sensitive communities can carry heightened obligations. Public-sector records and Indigenous data may also be subject to governing rights beyond an individual's click-through consent. Participation should be understandable, voluntary where required, and supported by meaningful dispute procedures. A contribution record must not expose a person's private identity simply to prove that money was allocated.

This model does not give a person automatic royalties from every past use of data associated with them. Economic participation depends on enforceable arrangements and applicable law. The social opportunity is to provide clearer, auditable ways to negotiate participation in future authorized uses of valuable contributions.

Part III. The seven transformations as a connected economic system

17. Transformation One: from AI assistance to distributed invention

Artificial intelligence is often sold as a tool that answers questions or accelerates existing work. NIKI and the FarmNode network take a different approach by connecting the generation of candidate inventions to testing, correction, persistent knowledge and later economic use. One laboratory can investigate a problem while another independently challenges the result. A qualified capability can become useful far beyond the organization in which it originated.

The large economic possibility is cumulative productivity: one validated invention reduces the effort needed for the next. The important quantity is not the number of AI messages produced, but the number of unfamiliar, useful methods independently established relative to the full cost of getting them there.

18. Transformation Two: from buying entire applications to buying useful capabilities

A whole software application is not always the smallest meaningful commercial product. DataChips make individual algorithms and functions identifiable. Circuits make them combinable. Circuit Genomes make their working arrangements reconstructable. The Composer and Universal SDK make this practical for developers and other software.

An independent inventor can potentially supply one exceptional component to several businesses instead of attempting to become the owner of a complete app marketplace. Customers can purchase the function they actually need, while permissions and technical evidence remain attached to the correct version.

19. Transformation Three: from easily copied bytes to portable governed rights

Digital files can be copied almost instantly; legal permissions cannot be inferred from a copy. ENTITY, ADAM, the Transaction Fabric and Verifiable Reality connect the identity of an asset with the evidence of its origin, the terms of lawful use, a governed action and an auditable outcome. The result is not simply moving data between servers. It is a framework in which properly established economic rights remain understandable across compatible organizations and technologies.

The potential significance extends to AI agents, which can recommend or request business actions without being given unlimited control of business records and accounts.

20. Transformation Four: from accumulated storage to accumulated knowledge capital

A file stored ten years ago may still be useful, but its value may be difficult to recover if nobody remembers its origins, assumptions or relationship to later work. BTDU links authoritative records to experiments, failures, corrections, meanings and permitted downstream uses. A retained unsuccessful experiment can be economically valuable if it prevents repeated dead ends. A verified improvement can become a reliable starting point for a new invention.

The transformation is not about turning storage space into money. It is about making verified knowledge more reusable and helping its economic usefulness compound across authorized participants.

21. Transformation Five: from company-only investment stories to asset-first enterprise formation

Today an investor commonly evaluates a whole company and all its operations. Bit Trade Global adds the possibility of showing exactly which data and intellectual assets underpin part of the company's productive potential. A documented, rights-clear asset with paying customers may support the formation or financing of a business structured around it. The appropriate securities still belong to a lawful issuer, not to a file hash or technical capability by default.

If this becomes a trusted institutional practice, more ideas from researchers, small firms and independent inventors could find commercial customers and appropriate financing without first becoming large corporations. The change is in the starting point for capital formation, not an exemption from company or securities law.

22. Transformation Six: from invisible contributors to identifiable participants

Contributors may currently have little insight into where their authorized input was used or how much income it helped earn. Through ENTITY's lineage, consent and entitlement records and the financial allocation infrastructure, a contributor can have a documented relationship to permitted downstream uses and actual contract proceeds. Cooperatives and collectively controlled assets may negotiate more effectively when their rights and allocation terms are clear.

The transformation is social as well as technical. It can move some bargaining power and appropriate economic returns closer to those who legitimately contribute, without assuming every person owns every dataset that mentions them.

23. Transformation Seven: from separate industries to one productive feedback cycle

Research creates knowledge. Knowledge becomes working capabilities. Rights and evidence allow those capabilities to reach customers. Actual customer payments support operating activity and legitimate contributor participation. Lawfully arranged investment may finance additional development. New funded research can produce new useful knowledge, beginning another cycle.

That is the central proposition of the Knowledge Capital Economy: **a connected system in which invention, ownership, commercial demand, finance and further invention can reinforce one another**. It is not a scheme for automatic or infinite financial growth. Research can fail; costs are real; investments can lose value. The promise is a more accountable and accessible path between productive intellectual work and economic opportunity.

Part IV. One complete example: following an invention through every system

24. A robotics invention, from the first problem to a new research project

The following is an explanatory end-to-end example of the fully interconnected architecture, **not a claim that BTG has completed an external sale, securities issuance or all of the steps as one production transaction**. The purpose is to show a non-technical reader exactly how the pieces are designed to work together.

1. The problem is defined. A manufacturing plant has a recurring fault that wastes production time. Its engineering team states the problem in measurable terms and identifies which observations it is permitted to provide to researchers.

2. The data is identified and restricted. ENTITY records the legitimate data controller, the source agreement and permitted research uses. Private machine and customer information remains protected. BTDU receives the authorized source record or governed reference, including its version and provenance.

3. A FarmNode opens the research project. A virtual robotics laboratory is assigned the permitted tools, data references, computing resources, limits and evaluation procedure. The research mandate identifies the human sponsor and the authority under which the study may proceed.

4. NIKI investigates. It reads relevant permitted knowledge and generates candidate algorithms. It does not start by treating the problem as solved; each candidate is a proposition that must be tested.

5. A candidate fails. An initial formula produces errors on a withheld test. That version and the negative results are retained. Researchers can later see why the design did not work.

6. A revised solution is created. NIKI produces new source code. The exact revised bytes are identified and frozen for qualification. Any later modification becomes a separately identified version.

7. Restricted execution takes place. The software runs with a defined memory and execution envelope. A testing service checks it against expected outcomes and negative security scenarios. A code source cannot be permitted to access unrelated company files merely because it was generated by AI.

8. Another evaluator tests it. The independent evaluator examines an unchanged version under its own conditions. The evidence reports both what passed and what limitations remain. A passing laboratory test is not yet proof of value at every factory.

9. BTDU preserves the research. It retains the source reference, experiment setup, inputs, outputs, failures, successful correction and relationships to earlier knowledge. The original source can be reconstructed when permitted.

10. The invention becomes a DataChip. Its precise function is packaged with a defined input, output, version, dependencies, supported execution environment, relevant qualification evidence and permitted distribution status.

11. DataChips form a Circuit. The new fault algorithm is connected to other functions: input validation, environmental adjustment, output classification and recording. The CircuitRuntime executes the actual connected computations rather than merely displaying boxes.

12. A Circuit Genome retains the design. The exact components and their connections are preserved in a reconstructable versioned architecture. A future revision can be compared with this working baseline rather than silently replacing it.

13. Another FarmNode reproduces the result. A separate research organization receives a rights-permitted capsule and conducts its own experiment. It signs the results it actually observed. A third lab may challenge unusual cases. Their identities and results remain distinct from the originating farm.

14. ENTITY establishes the rights chain. Evidence of authorship, employer or contributor agreements, source-data restrictions, assignment documents and licences are assembled into the relevant records. An asset is not treated as marketable unless the lawful commercialization rights exist.

15. Verifiable Reality labels the claims. A model-generated suggestion, an internal test, an independent technical result, a contractual entitlement and a bank receipt appear as different kinds of evidence. No one kind silently impersonates another.

16. A commercial asset is described. Bit Trade Global presents a non-sensitive explanation of the algorithm's use, limitations, supported operating environments, permitted licences and actual customer evidence. The underlying protected information need not be posted publicly.

17. A customer selects a licence. A manufacturer chooses the form of use it needs: perhaps the ability to run the calculation for one year, without permission to redistribute the source. Its contract and eligibility are verified.

18. ADAM authorizes the actual operation. A customer request is checked against current permissions and the expected version. Authorized execution proceeds. Unauthorized source export or use after expiry is blocked.

19. The Universal Transaction Fabric records the event. Signed event references connect the contracted rights, permitted execution, version used and delivery record, without pretending the record itself is payment.

20. A bank or payment service confirms actual money. The confirmed receipt is recorded separately from any quotation, promised amount or hypothetical valuation.

21. Legitimate contributors receive their share. If applicable agreements require revenue participation, the financial waterfall calculates entitlements from the available receipt. It handles fees, withholding, disputes and payment status, while preventing the same receipt from being allocated more than once.

22. The asset's business case becomes clearer. Customer renewals, verified savings and operating costs provide information for an evidence-based valuation. Market demand can inform the research agenda, but it is not proof of scientific correctness.

23. A lawful issuer may seek finance. A company holding the relevant commercialization rights may prepare an eligible offering with legal advice, disclosures and investor protections. Buying those securities gives the investor their specified financial rights, not automatic access to the source data. A live secondary market requires its own separate authorization.

24. New research is funded. Real available company resources or correctly authorized financing support a new research mandate. NIKI and other FarmNodes investigate the next problem. BTDU preserves the relationship between the earlier invention and the new work. The economic and scientific records show what was genuinely produced with the resources committed.

Seen as a complete chain, the system is not just 'AI plus a marketplace.' It is a new kind of working relationship between research, executable software, rights, independent evidence, real customers, contributor contracts and legally governed capital.

Part V. What the completed architecture means for industry and everyday life

25. Forestry, manufacturing and industrial maintenance

British Columbia's resource industries rely on machinery working reliably in demanding conditions. A mill, vehicle fleet, mechanical workshop or equipment manufacturer can collect lawful operating measurements that reveal faults, energy use and unexpected downtime. The challenge is not merely finding a mathematical pattern but making a solution dependable, transferable and usable by different owners of equipment.

NIKI and participating research FarmNodes can investigate algorithms using authorized records. BTDU preserves the experiment history. A DataChip packages a particular calculation; a Circuit combines it with input validation and equipment-specific rules. ENTITY limits how another manufacturer can use the resulting capability and protects source licences. A buyer may use a qualified method remotely or inside equipment without receiving every record in the originating dataset.

For a small engineering company, this structure could create a business based on one very useful industrial capability. For a larger enterprise, it may reduce the time needed to evaluate outside innovations because the source, rights and evidence are documented. For researchers, it offers the possibility that a narrow technical breakthrough can reach multiple sectors through different authorized licences.

These are industrial opportunities, not promises that any specific mill or manufacturer has agreed to purchase a BTG product. Real benefits would be measured through lower maintenance expense, fewer faults, longer equipment life, energy savings or other customer-verified outcomes.

26. Agriculture, water, ecosystems and rural research

Farmers and agricultural researchers work with weather, soil, water, machinery, growth conditions and market constraints. Local conditions matter, yet research performed in one location can sometimes contain methods that are useful somewhere else. The important distinction is between sharing a method and exposing someone else's raw private information.

A research FarmNode could investigate a moisture-use model based on participating farms' permitted observations. A DataChip might calculate a water-use risk score or flag an equipment anomaly. A Circuit could combine local measurements with a seasonal weather model and a grower's own operating constraints. The Genome would preserve how the parts interact, allowing later versions to be compared. ENTITY would record which source data may be used for research, trained models or commercial services and which may not.

The BTDU knowledge record would include failed trials and the conditions under which an observed improvement occurred. This is especially important in agriculture, where a method that worked on one crop, soil type or season may be unreliable in another. More information stored is not inherently better information; verified relevance and lawful reuse are what matter.

For the Boundary region, the practical opportunity could include digital agronomy skills, sensor research, collaboration with land-based businesses and specialized rural software development. Engagement with Indigenous knowledge and community-governed ecological information must respect the appropriate collective authority, not assume all research information can be freely bought or uploaded.

27. Search and rescue, emergency management and infrastructure

Search and rescue and emergency management depend on timely knowledge of terrain, water, roads, hazards, weather and field conditions. Tools that help interpret permitted information can support planning and situational awareness, especially in places where mobile networks are unreliable.

A BTG-based application might contain separately qualified DataChips for terrain interpretation, route distance, location uncertainty and weather comparison. A Circuit can combine those capabilities into an approved display or decision aid. BTDU preserves data sources and what they were measured to mean. ENTITY enforces restrictions on sensitive incident information. ADAM prevents an unauthorized research agent from altering incident records or issuing operational commands. A device can use offline source packs when permitted and available.

The human chain of command remains essential. An AI suggestion is not a direction to deploy rescue personnel into an unsafe area. Public safety use requires professional evaluation, suitable operating procedures, privacy protection and evidence appropriate to the risk.

The economic potential here is partly direct software utility and partly public value: better preparation, clearer information and potentially reduced duplication of research across agencies. Public benefit must be measured through appropriate field testing rather than inferred from model-generated maps or simulation scores.

28. Education, accessible learning and scientific training

The same software architecture can support learning without requiring students to use the financial exchange. A reading or science application may use a DataChip that adjusts the difficulty of a lesson, another that evaluates a mathematical answer, and another that controls a real programmable robotics activity. A Circuit can combine these functions into a personalized but accountable learning sequence. BTDU can preserve curriculum sources and measured learning progress while maintaining appropriate controls over children's information.

Accessible learning is not simply using AI to rewrite a textbook. Students with different learning needs may require specific accommodations, independent progression in literacy, mathematics, science and programming, and evidence that an adaptation actually helps. The technology must be evaluated for effectiveness and safety in appropriate settings. Families, educators and institutions should be able to understand what information is collected, who sees it and how long it is retained.

For colleges, universities and research institutions, FarmNodes could also support supervised student research and reproducibility training. Students can learn how an invention moves from hypothesis to software to independent testing while understanding that a computer-generated explanation and a scientific result are not the same thing.

29. Health, research ethics and privacy-sensitive computing

Health and life-science research can have substantial value, but it also demands particularly strong privacy, clinical and ethical boundaries. BTG's architecture could support arrangements where a qualified calculation runs against permitted information inside an institution, returning only authorized outputs rather than copying identifiable patient records into an open exchange.

A hospital research team may evaluate a method for detecting a pattern in de-identified measurements under a research agreement. The source data remains in controlled storage. A DataChip represents the relevant calculation. ENTITY records the permitted project and users. ADAM enforces the approved purpose, and BTDU retains appropriate technical evidence and lineage. A commercial partner does not thereby receive the right to view every patient record or use the method as a clinical diagnostic.

Clinical usefulness, regulatory status, informed consent and institutional review are separate questions. The presence of a Rights Passport or passing software tests does not establish clinical safety. The system's value is the possibility of precise, auditable cooperation without a rule that all knowledge must move into one central company database.

30. Universities, inventors, small businesses and new companies

Universities and smaller research teams often generate valuable specialist methods long before there is a customer-facing product. Turning them into complete applications, licensing them to industry and negotiating legal rights can be expensive. A qualified DataChip or Circuit Genome offers a more precise unit of technology transfer. A research method can be described with reproducible technical evidence and offered under a carefully defined use agreement.

A small organization may be able to commercialize a narrow capability without financing an entire platform. A manufacturer can assess that capability against its existing systems. An intellectual-property adviser can examine rights and attribution. A bank or investor can evaluate the operating business built around genuine demand. The technology does not remove legal and business costs, but it can make the subject of a transaction clearer.

This creates opportunities for new companies to **form because a useful intellectual asset has been identified**, rather than requiring every discovery to be enclosed within an existing large corporation before economic value can be recognized. That potential is especially interesting for regions where talent exists but venture capital, large laboratories and specialist commercialization services are less concentrated.

Part VI. The economic opportunity and how BTG could earn revenue

31. An existing large data economy provides context, not a promise of BTG revenue

Companies and institutions already spend substantial amounts creating, buying, managing and monetizing information. Industry researchers estimate a North American data-broker market measured in the tens or hundreds of billions of U.S. dollars, depending on the definition and period. These figures show the scale of existing commercial activity around data. They do **not** measure the turnover of a new public market for independent data-investment securities, and they should never be presented as revenue already available to BTG.

BTG's economic research dossier identifies a Fortune Business Insights estimate of **US$304.48 billion** for the 2025 global data-broker market and a reported North American share of approximately **41%**. This implies an estimated **US$124.84 billion** for North America under that research firm's category definitions. Because the regional share is rounded, this number is a derived approximation. Grand View Research separately estimated approximately **US$114.55 billion** of North American data-broker industry revenue in 2024. The methodologies, dates and categories differ; the figures are alternatives for context, not amounts to add together. [1] [2]

The case that information itself can be treated as productive capital is supported by a different set of studies. Statistics Canada estimated roughly **C$57 billion** in Canadian business investment in data assets in 2019, including internally created and purchased data assets. A March 2026 experimental working paper from the U.S. Bureau of Economic Analysis estimated **US$875 billion** in internally created U.S. data and database investment for 2024 and a related current-cost net capital stock of approximately **US$1.966 trillion**. These are statistical estimates of investment and accumulated capital, not sales through a BTG exchange and not official valuations of particular software algorithms. [3] [4]

The economic opportunity is larger than one market-research headline and more specific than a generic claim that 'data is worth trillions.' Productive information can be an input into manufacturing, transportation, agriculture, scientific research, digital services and many other industries. BTG's architecture addresses how identifiable portions of that useful knowledge may be verified, reused, commercialized and in suitable legal circumstances financed.

32. Revenue sources for a connected BTG business

The company could conduct several legitimate types of economic activity within the architecture. Businesses may pay for enterprise software and integration tools, qualified DataChip execution, research infrastructure, licensed BTG-owned inventions, metered commercial rights transactions, professional technical qualification, hosted or local computation, documentation services and specialized market information. An authorized investment-market operator, if established through the necessary legal arrangements, may have separate permitted fees. These streams have different costs and regulatory requirements.

BTG-owned assets such as a robotics algorithm or other legitimately controlled intellectual property may produce direct licensing income. Third-party inventors can also use the infrastructure without automatically giving BTG ownership of their work. Infrastructure revenue and asset ownership revenue must be kept distinct so that the economic claims of participants remain understandable.

There is an important difference between the size of the economy the platform supports and the amount the platform earns. A marketplace handling C$1 billion of permitted transactions is not automatically a business with C$1 billion of revenue. A company may earn a fee on a transaction, carry costs to process it, and retain a much smaller amount. The value of companies issuing securities is likewise not the same as the revenue of the exchange operator.

The dossier includes a proposed **2.9% plus C$0.30** commercial transaction charge. That is a design proposal, not established net profit. For very small machine-to-machine transactions, a fixed thirty-cent charge can be disproportionate. If an external payment processor also charges a comparable amount, most or all of the apparent gross fee could be consumed before computing, security, support and compliance expenses. Batching, prepaid balances, enterprise contracts or differentiated fee arrangements may be appropriate in a mature service, subject to actual terms and accounting.

The strongest commercial case for BTG would be demonstrated by unrelated customers who keep using valuable, rights-clear capabilities at prices that cover the entire cost of providing them. A real payment is stronger evidence of demand than an internal estimate, a simulated order or a public market forecast.

33. What could make knowledge-compounding economically powerful

If a laboratory develops one useful capability, it has one potential product. If that capability becomes a reliable component used in many Circuits, it may support multiple products. If later researchers can preserve and build upon its evidence, it may reduce repeated research cost. If customers pay for lawful uses, part of the resulting economic surplus may finance more experiments. These effects create a possible feedback loop.

The compounding effect is **conditional**. Reusing a component must save more resources than it costs to qualify, store, support and license it. A network of farms must produce real new utility, not just recycle the same code under new names. Customers must have actual problems they are willing to pay to solve. Financial valuations must not be inflated by repeated internal transactions among related parties. The value of an unsuccessful experiment lies in genuine learning, not a fictitious dollar amount assigned to failure.

The complete system creates a way to observe the cycle rather than merely describe it. An operator can measure qualified discoveries, independent reproductions, real customer use, actual receipts, net costs, contributor allocations and the incremental output of reinvested research. When those records are sound, the economic significance can be examined on its merits.

Part VII. Grand Forks, British Columbia and the wider public interest

34. Why a global knowledge economy can originate in Grand Forks

Grand Forks is not normally thought of as a global financial or artificial-intelligence centre. Its importance to this project is precisely that location should not determine who can create useful intellectual work. Many computational research and software activities can be performed from smaller centres when the necessary skills, telecommunications, computing capacity, partnerships and governance are available.

For the Boundary region, the relevant opportunity is not a promise that one technology company will transform the local economy overnight. It is the possibility of developing a specialized regional capability in software engineering, AI evaluation, applied research, industrial data systems, rights and provenance services, secure infrastructure and professional commercialization support.

A local research company can work on problems that matter in a rural region - equipment reliability, forestry operations, agriculture, water, accessibility, emergency preparedness and technical education - while designing solutions for customers across Canada and internationally. Some of the value could remain local through employment, training, contracting and the development of additional firms. The size of those benefits depends on actual investment, customers, skills development and long-term business performance.

A community should also ask about infrastructure and risk: internet reliability; access to qualified workers; data-centre energy and water requirements where relevant; housing and business costs; cybersecurity; the value of training partnerships; and whether public expenditure is justified by independently measured local benefits.

35. Opportunities for British Columbia and Canada

British Columbia has internationally relevant industries in natural resources, clean technology, digital services, research and software. Canada has research universities, financial institutions and established regulatory structures. BTG's work touches these areas through three different kinds of partnership.

Applied research partnerships can evaluate a specific invention method or DataChip against a useful outside baseline. **Commercial partnerships** can licence tested capabilities and measure real operational gains. **Institutional and policy discussions** can examine what is required for privacy-preserving intellectual-asset finance, contributor rights and regulated market integrity. Those tracks should be assessed separately rather than presented as a request that one public institution endorse the whole architecture.

Canadian research institutions could examine whether independent FarmNode replication reduces time and expense, or whether BTDU improves the ability to reconstruct a complicated experiment. Manufacturers could compare a qualified algorithm with the existing method they already use. Indigenous governments and communities could determine whether any cooperation supports their own data-governance priorities; the ability to participate never implies an obligation to share or sell information.

A provincial or federal economic-development conversation can focus on productivity, technical training, rural participation, responsible commercialization and transparent governance. It should not be confused with an application for permission to operate an exchange, which falls under the appropriate securities and other legal authorities.

36. How public bodies could participate without surrendering their responsibilities

A city, regional district, First Nation, province, university or other public institution has its own mandate, privacy obligations and procurement rules. Participation in technical review or a research challenge does not require endorsement of BTG's financial platform. A public agency could, for example, publish an impartial problem statement about safe equipment monitoring, require open and fair participation where applicable, commission qualified independent testing, or consult relevant experts about privacy and evidence standards.

Access to records and data must follow the public body's governing law and institutional policies. The fact that a local technology company can make useful use of information does not automatically mean the public body may give it confidential records or exclusive commercial rights. Public funding and purchasing decisions must serve documented public purposes under normal processes.

Where a private-sector founder has a family, financial or other personal relationship with a public office-holder, engagement requires particularly clear separation between the public interest and private business negotiations. Relevant interests should be disclosed, conflicts assessed under applicable law, and any required recusals or impartial decision-making processes observed. Personal acquaintance with a mayor, director, councillor, MLA, MP, minister or other official does not confer official approval or exemption from procurement or securities rules.

This protects both the public and the credibility of BTG. A scientific result is more convincing when outsiders can reproduce it; likewise, a public partnership is more convincing when the institution has made its own independent decision under a fair and documented process.

Part VIII. Trust, regulation and practical boundaries

37. Privacy, intellectual property and Indigenous data governance

The system's engineering separates public discovery of an asset from private access to its source information. A catalogue may say that a qualified algorithm exists and is available for a specified use. That is not permission to publish the data used to develop it, identify a person mentioned in a source record or transfer a community's confidential knowledge. Sensitive data may have to remain in a controlled institution, with computations brought to it under strict conditions or the entire use refused.

Canadian privacy and public-sector information laws apply according to the organization and activity. In British Columbia, private organizations and public bodies may be governed by different statutes. Institutions should identify the correct legal basis, permitted purposes, affected individuals' rights, retention and security obligations before granting access. Labour agreements, intellectual-property contracts, research ethics and cross-border export restrictions may also matter.

Indigenous data sovereignty is not exhausted by a checkbox or individual consent screen. Indigenous governments and communities may have collective governance principles and legitimate authority over information about their peoples, territories and knowledge. Any system involving such information must be developed through appropriate relationships, defined authority and community-controlled limits. Sometimes the correct decision is that the information should not enter a marketplace.

38. Why legal rights cannot be generated by software

An algorithm can be tested and reproduced. A file's bytes can be signed and compared. A public catalogue can show an issuer's statement about an invention. None of those actions automatically gives the issuer the legal right to sell it. Legal title may depend on employment, assignment, confidentiality, copyright, patents, licences, privacy and the rights of contributing organizations.

Likewise, a database entry saying someone owns 1,000 investment units is not enough to establish that the units were lawfully issued shares. The governing corporate documents, statutory or authorized registers, investor eligibility, source of funds and applicable securities regime matter. A wallet should reflect legal reality rather than attempt to replace it.

This is not a weakness in the concept; it is the architectural reason ENTITY, ADAM, Verifiable Reality and the separate financial registries exist. The system is engineered to make such boundaries explicit so autonomous software cannot create legal or financial authority merely by calling a convenient API.

39. Why a commercial data marketplace and a securities exchange are different

A customer paying for permission to run a DataChip normally enters a commercial agreement for utility. An investor buying shares in a company accepts different risks and obtains different legal rights. Securities are subject to particular rules about issuance, disclosure, distribution, advice, dealing, resale, custody and marketplace activity.

In Canada, the B.C. Securities Commission administers relevant provincial securities law. Federal and national frameworks, including Canadian Securities Administrators instruments and the rules of CIRO where applicable, shape the requirements for dealers and markets. An alternative trading system and a recognized stock exchange are regulated institutional activities. A consultation letter or approved private primary securities offering does not itself authorize public secondary-market trading. [5] [6]

BTG's engineering dossier therefore separates a commercial marketplace for permitted data and software uses from any live investment-market function. The investment layer can be specified, simulated and independently

examined while the required legal operating arrangements are pursued through the appropriate channels. That is a more credible foundation than assuming that describing a share as a 'bit' avoids securities law.

40. Security and responsible autonomy

An agent that writes and runs software must be treated differently from a simple reporting tool. Generated code can contain errors, consume excessive resources or attempt access outside its assigned project. Research execution therefore belongs inside a bounded environment with recorded limits. The relevant operating system or virtualization layer, not an AI's own promise, must enforce critical restrictions.

Across the broader ecosystem, identity verification, short-lived permissions, access limits, encrypted communications, recorded approvals, versioned software, signed evidence and independent security testing protect users and organizations. Cash and securities records require exact financial reconciliation. A failure should lead to a defined halt, quarantine or correction instead of an invisible override.

No technical design can guarantee the absence of all security incidents. The appropriate claim is that the architecture specifies controlled permissions, separation of authority, observable results and independently reproducible evidence. Those controls must be tested in the environments in which the system is actually used.

Part IX. The demonstrated engineering foundation and the complete system being built

41. The October 2026 research evidence

BTG's technical record includes actual local research, code execution, knowledge reconstruction, Circuit execution and encrypted research-reference delivery. Those records are more substantive than a proposal consisting only of diagrams and written specifications. They also concern specific, bounded experiments, not yet a verified global commercial economy.

One research challenge required an algorithm to decide whether at least three of four encoded sensors reported the same valid reading. There were 65,536 possible input combinations. NIKI's first recorded generated C expression gave incorrect answers for 57 combinations, including cases in which a missing sensor was mistakenly counted as an agreeing measurement. The failure was preserved and the repair process produced a corrected source version with its own recorded provenance.

The corrected source was compiled to a 240-byte WebAssembly module. The recorded Windows Wasmtime execution returned correct results for all 65,536 inputs. A later Debian WSL2 execution using Bubblewrap and cgroup controls also returned 65,536 correct results. The recorded compiled binary hashes matched. The Linux environment ran on the same BTG workstation, so this is a real second operating-environment test, not a claim of reproduction by an independently controlled external institution.

The associated BTDU and ADAM records preserved the original research material and allowed local cold reconstruction. The Recursive Fabric runtime initially executed an evidence-processing Circuit Genome connected with the research chain. Later engineering advanced the demonstration: a recorded local receipt shows the actual repaired NIKI sensor algorithm operating as a typed executable node in a two-node Circuit, followed by an evidence-digest operation. The complete result matched the frozen oracle across all 65,536 inputs. Subsequent internal engineering records describe an immutable DataChip descriptor retained and reconstructed through BTDU, along with targeted and negative-test results. These findings form part of the documented internal engineering record and are distinct from external independent certification.

The record also includes three encrypted research-reference deliveries to **three local Windows worker processes** with durable acknowledgments and duplicate handling. Those workers were not three separately controlled physical research institutions, and the deliveries did not create legal commercial rights. An earlier downstream suite recorded 312 local regression tests passing after a process-launch correction; a later engineering report records 356 Windows backend regression tests passing as the integration progressed. These are software tests within their stated scopes, not counts of inventions or market transactions.

The relevant historical artifacts are held in the BTG qualification workspaces, including source and binary hashes, execution receipts and research delivery files. These provide a basis for future external reproduction. These development records do not establish independent institutional certification, owner-authorized production DataChip issuance, three separately enrolled physical hosts, a completed customer sale or authorized securities trading. Each is a distinct matter requiring its appropriate evidence and authority.

Why this real engineering example matters to a public reader

It demonstrates the *kind* of continuity the architecture is meant to preserve. A candidate can be generated, found wrong, repaired, executed as working software, reconstructed from retained knowledge, connected into a larger working Circuit and transferred as a research reference - while each stage has its own record. The more important later economic and institutional claims become credible only when independent organizations and paying customers complete their corresponding stages.

42. How the engineering architecture is organized

The seven transformations are being engineered as one interoperable system, not as seven unrelated smartphone apps. Each major service has a distinct responsibility and a documented interface to the services around it.

The discovery side contains NIKI, its research mandates and the computational FarmNodes. The knowledge side contains BTDU's authoritative payloads, atoms, bonds and reconstruction. The executable side contains DataChips, AI Chips, Circuits, Genomes, the Composer and supported runtimes. The authority side contains ENTITY, ADAM, the Transaction Fabric and Verifiable Reality. The economic side contains asset registries, commercial licence workflows, contributor agreements, valuation evidence, fiat payment reconciliation and distinct legally governed investment records.

This arrangement avoids requiring a single enormous database to act simultaneously as a research notebook, source-code repository, privacy office, patent registry, bank and securities registrar. Common identifiers and verifiable records allow systems to communicate, while the proper authority for each fact remains identifiable.

The engineering package also contains typed data contracts, API definitions, system-boundary requirements, adversarial tests, migration rules and an inter-system demonstration sequence. Those documents direct engineering; their mere presence does not establish that every described capability is functioning. Continued build, test and external review provide the evidence appropriate to each use.

Part X. What this infrastructure could become

43. A global distributed research network

At one scale, BTG could be a valuable engineering company whose services help customers develop and use trustworthy software components. At a larger scale, independently operated FarmNodes could form a network of virtual research laboratories with distinct owners, specializations and technical responsibilities. Each lab could investigate unfamiliar problems, reproduce permitted discoveries and contribute qualified capabilities to other participants.

The strongest form of this network would allow a small laboratory in rural British Columbia to solve part of a problem for an overseas manufacturer, while a university independently verifies the method and another developer builds a compatible application. The originating parties would not have to surrender every source file, research asset or contractual right merely to participate. This is a form of collaboration that existing cloud and scientific systems support in fragments; the intended contribution is the connected evidence, capability and rights record.

A global research network could be important in manufacturing, energy, agriculture, materials, scientific software and many specialized fields. Its economic significance would depend on how often its results are independently useful, not on the number of virtual nodes shown on a network map.

44. A new category of software and intellectual-asset businesses

Applications are familiar products, but the capabilities inside them may become more visible economic units. A developer could specialize in creating qualified DataChips for a narrow industrial task. Another company could assemble several into a specialized Circuit. A third could maintain a Circuit Genome that can be adapted to different equipment. A research institution could licence an independently verified method. Each participant could have an identifiable relationship to the intellectual assets it legitimately controls.

This could broaden the forms of software entrepreneurship available to a small team. It could also change procurement: instead of buying a complete, expensive software suite merely to obtain one capability, customers might contract for a defined function and bring it into their existing applications. The engineering burden would remain significant, especially for safety-sensitive systems, but reusable evidence and rights contracts could reduce duplication.

45. A market for knowledge utility, not merely copies of data

One possible outcome is an economy in which the most important commercial product is the *benefit produced by a knowledge asset*, not a copy of the underlying file. A customer may obtain a fault prediction, an optimization result, a verified research method or an authorized computation without acquiring the training material or intellectual-property title. That is particularly valuable where privacy, trade secrets or territorial rules limit the movement of data.

The result would be a richer market of narrowly defined permissions and useful outputs. It would still coexist with ordinary software licences, open-source software, patent licences and conventional data markets. BTG's value would have to come from better integration, reliable evidence, useful rights portability, customer experience and operating economics, not from claiming that no alternatives exist.

46. Asset-first financing and the creation of new companies

The most ambitious economic change is that productive knowledge might more frequently become a documented starting point for company formation and investment. An identified dataset or invention with legitimate rights and credible buyer demand could support a specialized operating business. That business might finance further engineering, servicing, compliance, marketing and research through appropriate commercial or investment arrangements.

For investors, the advantage would be clearer visibility into the actual intellectual assets underpinning part of a company's business. For inventors, the advantage could be an additional route from narrow technical discovery to a financed enterprise. For society, the advantage could be more independent innovators reaching customers before their work becomes absorbed into a small number of large platforms.

A future mature institution of this kind could be economically significant, but its size must be discussed carefully. **Total value of listed assets, market turnover, issuer capitalization, customer productivity gains and BTG's own revenue are five different measurements.** They cannot be added into a single figure or presented as one

company's worth. A widely adopted infrastructure may support enormous economic activity while capturing a relatively modest fee for operating it.

47. Broader access to participation in technological wealth

A credible contributor economy could allow research cooperatives, small inventors, lawful data custodians and independent teams to negotiate shares of future revenue based on actual agreements. It could help make the origin of important intellectual assets more visible. It could preserve competing ownership models rather than requiring every participant to transfer all rights to a central exchange operator.

This outcome would depend on reasonable contract terms, affordable technical participation, dispute procedures, meaningful privacy choices and competition. A system in which only a few large buyers can acquire exclusive control over useful knowledge may reproduce the concentration it seeks to address. Portability, clear governance and actual contributor choice are therefore not decorative features; they affect whether the resulting economy spreads opportunities or narrows them.

48. Three possible scales of long-term economic impact

A strong Canadian technology business. BTG operates a reliable invention, software-component and rights platform serving research teams, manufacturers and professional software developers. Commercial revenue comes from customers who independently value particular capabilities. This would be meaningful even without a global securities-market operation.

An international intellectual-capability ecosystem. Outside laboratories and companies routinely create and reproduce DataChips, Circuits and Genomes; customers licence them under compatible terms; contributors receive appropriate contractual proceeds; and more institutions treat qualified knowledge as a distinct economic resource. This would represent a broader change in software and research commercialization.

Institutional knowledge-capital infrastructure. Asset-specific economic evidence becomes sufficiently standardized and trusted that a wider range of businesses, financiers and qualified issuers use it to assess, govern and fund the productive potential of intellectual assets. In that scenario, the economic effects may extend to business formation, research finance, digital rights and the recognition of data as productive capital. This is the most far-reaching possibility, and it depends on institutional trust, regulation, real demand and independent adoption, not on software engineering alone.

These futures are descriptions of what the completed architecture could enable, not financial projections or guaranteed outcomes. The quality of the system will ultimately be judged by real-world usefulness, verifiable rights, cost-effective operation and credible public trust.

Part XI. Questions a reader may reasonably ask

49. Is NIKI simply another chatbot?

No. NIKI is being engineered as a research and invention system that generates executable candidates, investigates failures and connects to independent tests, retained knowledge and reusable capabilities. It may use language models as components, but a fluent answer is not treated as a verified invention. The difference becomes measurable through actual working outputs and reproducible results.

50. Are DataChips physical chips?

No. They are digital software capabilities with defined behaviour, interfaces, versions, technical evidence and rights. They may ultimately run on physical processors or devices, but the economic unit is the qualified capability rather than a piece of silicon.

51. Is a Circuit an electrical circuit?

No. It is a connected assembly of software capabilities. A Circuit defines how the output of one DataChip becomes the input of another and how the complete process is executed and verified.

52. Why are Genomes needed when ordinary software has versions?

Ordinary version control preserves changes in source files. A Circuit Genome adds the complete structure of an executable assembly: its components, relationships, interface contracts, dependencies and reconstruction plan. A Farm Genome does something similar for an entire research laboratory. They are not replacements for Git; they are higher-level descriptions of how working systems are formed and evolved.

53. Is BTDU just a database or a large language model?

BTDU includes structured knowledge and authoritative content, but its purpose is broader than a single ordinary database or trained model. It retains source payloads, meaningful relationships, evidence history, reconstruction information and controls over who may use a given source. It helps NIKI and other systems work with traceable knowledge rather than pretending that a generated summary is the original fact.

54. Does ENTITY decide who legally owns information?

ENTITY records claims, supporting documents, identities, contracts and permissions and helps verify them. Legal ownership still depends on valid underlying rights and the appropriate laws and authorities. A valid digital signature proves something about a signed record, not automatically the truth of the legal claim inside it.

55. Why does BTG need ADAM as well as ENTITY?

ENTITY describes the principal, asset and relevant rights. ADAM is the controlled decision process that decides whether a particular requested state change is authorized at that moment and records the result. Knowing that a person has permission is different from permitting a specific action on a specific version under current rules.

56. Are computational farms thousands of computers owned by BTG?

Not necessarily. A FarmNode is a logical laboratory with its own identity, tools, rights, data and operating policy. It may use physical hardware controlled by BTG or another organization. A large network could include independently operated farms, subject to real trust and commercial arrangements. One workstation running several workers is a useful local development test but not proof of several independent institutional partners.

57. Must everyone release their raw data or source code?

No. Asset metadata and evidence can be discoverable while the underlying source remains private. A customer can sometimes buy a permitted calculation or execution service without receiving confidential training data or source code. Rights, privacy and data governance determine which arrangements are allowed.

58. Does a DataChip automatically have a market price?

No. It may have potential value, but commercial worth depends on legitimate rights, tested utility, actual customer demand, competition and cost. A technical identifier or computer-generated valuation does not create a real buyer or an enforceable economic entitlement.

59. Does purchasing a licence make someone a shareholder?

No. A commercial licence confers specified use rights. A share is a legally issued financial security with distinct rights and obligations. Contributor royalties are different again. The BTG architecture intentionally keeps them separate.

60. Is Bit Trade Global a cryptocurrency exchange?

That is not the controlling design. Fiat currency and approved settlement channels are used for real payments. Digital tokens or signed receipts can provide evidence of contracts or transactions without becoming a speculative general-purpose coin. Securities activity remains subject to its own rules.

61. Why is this relevant to a small community?

The creation of software and verifiable digital research is not confined to major metropolitan laboratories. With skilled people, suitable infrastructure and outside partners, a rural region can contribute valuable expertise and capabilities to national or international customers. The effect on jobs and local prosperity would depend on actual business performance and training, not merely the location of the company's headquarters.

62. Has everything described in this paper been commercially launched?

No. The work is active production-oriented engineering with documented local component experiments and an extensive interconnected architecture. The complete independent multi-institution research-to-customer-to-investment cycle is not presented here as an already settled commercial and regulated transaction. That distinction allows readers to understand the system being built without mistaking engineering specifications for independent financial or legal authorization.

Conclusion. A different relationship between knowledge and economic opportunity

The central concept behind BTG is straightforward: useful knowledge should not have to remain invisible inside a corporation, a model, a repository or a private database. When its origin is legitimate, its behaviour is reproducible and its rights are clear, it can become a reusable capability, a product, a subject of commercial agreement and a documented contributor to a business's economic value.

NIKI creates and repairs candidate inventions. FarmNodes investigate and reproduce research. BTDU preserves knowledge and the evidence surrounding it. DataChips provide discrete executable capabilities. Circuits connect those capabilities, and Genomes preserve the structures that allow them to be rebuilt. ENTITY records identity, provenance, rights and relationships. ADAM governs which actions are allowed. The Universal Transaction Fabric connects the events. Verifiable Reality distinguishes what is merely claimed from what has actually been established. The Composer, Universal SDK and Command Centre help people build, deploy and support working systems. Bit Trade Global provides commercial exchange infrastructure and a distinct route toward legally constituted intellectual-asset finance. Contributor agreements allow legitimate participants to share in real proceeds.

The power of the design lies in the connection between these systems. Each has a different job. Together they could create a continuous relationship between research, knowledge, software, legally controlled rights, real customers, investment and further invention.

For Grand Forks and the Boundary region, this work creates an opportunity to participate in research-intensive industries without assuming that innovation must begin in a major metropolitan centre. For British Columbia and Canada, it presents questions about productivity, rural digital work, intellectual-property commercialization, trustworthy AI and responsible finance. For the wider world, the longer-term possibility is an economy in which qualified knowledge is more independently identifiable, usable and financeable, and in which legitimate contributors can participate more clearly in its value.

The Knowledge Capital Economy is being engineered to make that possibility concrete. Its credibility rests not on extravagant claims about AI or market capitalization, but on the actual functioning of the complete architecture, truthful evidence, useful commercial outcomes, protected rights and institutions willing to evaluate it on its merits.

Appendix A. Plain-English guide to the named systems

System / term	Meaning and function
NIKI CGSI-SI	An autonomous research and invention engine. Generate, test, repair and learn from executable candidate methods
FarmNode	A whole virtual research laboratory with its own identity and resources. Perform authorized research and independently reproduce permitted discoveries
Farm Genome	A reproducible description of a research laboratory. Recreate and evolve a laboratory's tools, processes and permissions
BTDU	Knowledge representation, source preservation and reconstruction infrastructure. Keep authoritative information and its meaningful, traceable relationships
Atom Universe	BTDU's network of fundamental symbolic units and typed relationships. Relate language, code, mathematics and research claims without erasing originals
DataChip	A versioned executable software capability. Perform one well-defined and potentially licensable function
AI Chip	A DataChip-like model-backed capability. Offer qualified AI behaviour with declared model limits and rights
Circuit	A working assembly of DataChips and other permitted operations. Execute a larger process in a defined order and dataflow
Circuit Genome	The versioned architecture and reconstruction recipe of a Circuit. Retain the composition, dependencies, lineage and testable evolution of the Circuit
DataChip Composer	Source-code editor and visual component assembly workspace. Build and connect real capabilities using code and visual design
Universal SDK	Cross-environment integration tools. Connect supported devices and applications to BTG capabilities
ENTITY	Persistent asset identities, provenance, rights and entitlements. Show what an asset is and which claims and permissions govern it
ADAM	Governed state-change engine. Check whether a specific action is authorized and record its outcome
Universal Transaction Fabric	Signed event and delivery infrastructure. Connect meaningful events between independent systems without double-counting
Verifiable Reality	Evidence classification and verification framework. Distinguish claim, experiment, independent review, legal right and settlement
BTG Command Centre	Authorized engineering and support infrastructure. Diagnose applications, devices and

System / term	Meaning and function
	integration problems with controlled access
Bit Trade Global	Commercial digital-asset rights marketplace and separate investment infrastructure. Support lawful asset commercialization and appropriately authorized financial instruments
DCO	An identified digital commercial or intellectual asset in BTG's architecture. Provide a distinct asset and associated rights reference
Rights Passport	A record of asset-specific permissions and restrictions. Explain who may perform which use and under what legal conditions
Wallet	User-facing view of entitlements, contracts, assets and permitted holdings. Show relevant records without replacing bank or legal registers

Appendix B. Selected sources and evidence

The economic and institutional evidence below provides context and precise definitions. Financial figures refer to their respective source dates and methodologies. BTG engineering references document local work, not independent third-party certification of every system.

Public economic and regulatory sources

[1] Fortune Business Insights, *Data Broker Market*. Estimated 2025 global revenue and regional share used for the stated derived North American estimate. https://www.fortunebusinessinsights.com/data-broker-market-116970

[2] Grand View Research, *North America Data Broker Market Size & Outlook*. Separate 2024 regional estimate. https://www.grandviewresearch.com/horizon/outlook/data-broker-market/north-america

[3] Statistics Canada, *Data, Intangible Capital and Economic Growth in Canada* (2025 paper reporting 2019 investment). https://www150.statcan.gc.ca/n1/pub/11f0019m/11f0019m2025003-eng.htm

[4] U.S. Bureau of Economic Analysis, *Toward Economic Statistics on Own-Account Data and Database Assets for 1997–2024*, Working Paper 2026-4. Experimental methodology, not official securities turnover. https://www.bea.gov/sites/default/files/papers/BEA-WP2026-4.pdf

[5] British Columbia Securities Commission, marketplaces and financial technology information. https://www.bcsc.bc.ca/industry/marketplaces-sros-market-infrastructure/marketplaces

[6] Canadian Investment Regulatory Organization, marketplace/dealer oversight information. https://www.ciro.ca/

[7] B.C. Community Charter, including conflict-of-interest provisions applicable to local elected officials. https://www.bclaws.gov.bc.ca/civix/document/id/complete/statreg/03026_00

[8] British Columbia, *Personal Information Protection Act*. https://www.bclaws.gov.bc.ca/civix/document/id/complete/statreg/03063_01

[9] British Columbia, *Freedom of Information and Protection of Privacy Act*. https://www.bclaws.gov.bc.ca/civix/document/id/complete/statreg/96165_00

[10] Bytecode Alliance, WebAssembly and WASI/component interoperability resources. https://bytecodealliance.org/ and https://component-model.bytecodealliance.org/

BTG engineering and architecture documents

[D1] BTG, *Data as the Asset: The Public Data Investment Economy*, October 19, 2026. Founder’s controlling asset-first economic thesis in the investment engineering dossier.

[D2] BTG, *All Seven Transformations: Integrated Board Directive* and seven-system architecture, source contracts, rights boundaries and research/economic qualification specifications, October 2026.

[D3] BTG, *Autonomous Invention Qualification* and NIKI source/Wasm research receipts, October 2026, including recorded sensor-oracle failures and repaired Windows/Linux local results.

[D4] BTG, *NIKI Typed WASM Circuit Receipt*, local Recursive Fabric execution evidence, October 2026; further DataChip descriptor and negative-test results reported by the ongoing engineering session.

[D5] BTG, *North America Data Economy Research and Source Proof*, October 2026; source-by-source distinctions between data broker revenue, data capital formation and financial securities-market activity.