Nuclear Power Industrial Processing Initiative

CLIMATE MOONSHOT PORTFOLIO

Billions to save trillions

Proposed $1 Billion / Five-Year R&D Program

31 INDUSTRIAL PROCESSING PLATFORMS

STANDARDS DEVELOPMENT

By Glenn Weinreb

Developed with assistance from ChatGPT

OPEN SOURCE

Original files are free to copy and modify under the CC BY 4.0 open-source license.

August 31, 2026 | v3d | Copyright © 2026 by Glenn Weinreb | www.APlanToSaveThePlanet.org

Executive Summary

Purpose. This proposed $1 billion, five-year R&D initiative develops standardized, factory-built industrial-processing platforms and the shared infrastructure needed to use low-carbon nuclear electricity and heat at large industrial sites. The reference architecture combines automated construction Chinese PWR and HTR heat/electricity assumptions with a portfolio of 31 repeatable industrial-processing platform families built on a 12 m-wide modular envelope, normally in 24 m increments up to a 96 m four-module platform, including low- and medium-temperature process heat, high-temperature heat, direct air capture, hydrogen and ammonia production, and supporting industrial utilities. It is an R&D and demonstration program, not a budget to build or operate a licensed nuclear-industrial complex.

Economic hypothesis. Factory fabrication, standardized interfaces, heat cascading, industrial heat pumps and mechanical-vapor recompression, modular transport, and repeated operation may lower the cost of replacing fossil-fired industrial heat and chemical production. PWR electricity and low/medium-temperature heat are emphasized where they are cheapest; HTR heat is reserved for processes that genuinely benefit from roughly 300–500°C service. Cost and climate benefits are treated as hypotheses to be measured, not guaranteed outcomes.

China strategy. China is used as the initial industrial-development and cost reference case because it combines an active nuclear build program, large chemical and heavy-industrial supply chains, major shipyards and ports, and manufacturing capacity capable of repeated production of large modules and process equipment. The program should first prove the 31-platform architecture, shared logistics, standards, heat interfaces, ammonia/hydrogen systems, and long-duration test methods under controlled Chinese assumptions and measure delivered product cost, reliability, energy use, emissions, logistics, and construction performance. If successful, sponsor-developed platform designs, machines, software, interface standards, test methods, and manufacturing packages can then be adapted to locally licensable nuclear and industrial systems in Europe, the United States, and elsewhere. China is therefore the first high-volume proving and scaling environment, not the exclusive destination for the technology.

Climate Money. The proposal assumes the funding source aims to maximize climate benefit per dollar spent rather than financial return on an investment. Public and philanthropic R&D funding can develop process-platform designs, transport and construction equipment, standards, software, prototype evidence, and qualification methods as public-good infrastructure, with sponsor-funded designs and data released for competitive reuse to the maximum extent lawful and technically appropriate. The $1 billion program would not finance commercial nuclear plants, commercial chemical plants, or a commercial ship fleet; subsequent public and private industrial capital would. Public R&D, grants, tax support, carbon prices, and other subsidies should remain transparent in the

economic accounting, while proprietary reactor-vendor information, controlled nuclear information, cybersecurity-sensitive material, and other legally restricted data remain protected.

Five-year work. The baseline technical program is $800 million with a centrally controlled $200 million risk reserve. It develops FEED/reference designs for all 31 platform families; substantial integrated pilots for roughly the highest-priority 8–12; hydrogen and ammonia systems; competing marine ammonia fuel-cell prototypes with long-duration testing; material-handling tunnels, storage and waterfront systems; heavy platform transport and ship/shore transfer; segmented electrified rail; machines for automated thermal-storage-tank construction; open interface standards and interoperability testing; and a non-nuclear prototype campus with PWR-like and HTR-like heat loops, full-temperature flow/corrosion tests, and endurance campaigns.

Success criterion. The objective is independently reproducible evidence on energy efficiency, process performance, availability, degradation, maintainability, safety, emissions, interface compatibility, transport and construction performance, cost, and CO_2 reduction. Success also requires open standards that allow multiple suppliers to connect equipment without bespoke redesign and decision-grade evidence for which platform families merit Phase II scale-up. The program does not guarantee specific electricity, heat, or industrial-product costs, and it does not require nuclear fuel. A successful China result should also identify which platform designs, standards, equipment, and test methods can be transferred economically to the United States, Europe, and other regions after local qualification, localization, and re-costing.

Climate Moonshot Portfolio. We define "Climate Moonshot" as conducting R&D to the extent required to solve the climate problem. This would require a portfolio of billion-dollar-scale R&D proposals, examples of which are available for free at www.APlanToSaveThePlanet.org/cmp. The document you are reading is one initiative within that portfolio. It uses a Common Framework establishing shared rules for governance, procurement, open-source and data rights, independent review, financial oversight, and stage-gated funding.

Program at a Glance

The program is a modular industrial-systems R&D effort. It develops process platforms and the shared logistics, standards, and decision evidence needed to determine which applications justify later scale-up.

Item	Program design
Total funding and duration	$1.0 billion over five years.
Mission	Develop standardized, factory-built industrial-processing systems that use low-carbon nuclear heat/electricity and can beat fossil alternatives on delivered cost and CO_2 reduction.
Reference energy cases	Automated China PWR for lowest-cost electricity and low/medium-temperature heat; HTR-PM for processes that materially benefit from roughly 300–500°C heat.
Regional sensitivities	United States and selected European cases re-price energy, labor, platform fabrication/localization, codes and qualification, port/logistics, financing, carbon/policy treatment, and fossil comparators using the same measured platform evidence.
Module architecture	12 m-wide modules in 24 m increments; 12 × 24 through 12 × 96 m platforms with open mechanical, thermal, electrical, material, rail, safety, and data interfaces.
Platform portfolio	31 platform-family reference designs; substantial integrated pilots for roughly the highest-priority 8–12; subsystem/test-article work for other viable families.
Major enabling systems	Hydrogen/ammonia, marine fuel cells, material handling, waterfront systems, heavy module transport, automated thermal-storage construction, and interoperability standards.
Test architecture	Non-nuclear prototype campus with PWR-like and HTR-PM-like heat loops, representative pressure/temperature/material conditions, corrosion rigs, and long-duration operation.
Budget	$800M planned technical program plus a centrally controlled $200M risk reserve.
Evidence and end product	Independent measurements of efficiency, throughput, reliability, degradation, safety, cost, interoperability, and CO_2 reduction; platform-by-platform recommendations for Phase II.

Reference Videos

Low-Cost Nuclear Power [#11]
https://www.youtube.com/watch?v=AIlbovU67wI

Automated Nuclear Power Construction [#12]
https://www.youtube.com/watch?v=afO0cy1l7Qo

How to Make $10 Trillion Dollars [#13]
https://www.youtube.com/watch?v=4gqmKGV1h5Y

Table of Contents

Executive Summary 2
Program at a Glance 4
Table of Contents 4
Nuclear-Integrated Industrial Processing Platforms 7
Relationship to the Nuclear Power Automated Site-Construction Initiative 11
Global Deployment Economics 13
Proposed $1B R&D Initiative 16
Program Architecture and $1 Billion Budget 19
Program Crosswalk 21
Stage Gates and Go/No-Go Criteria 23
Program-Wide Measurement and Acceptance Criteria 25
Development Phases 27
Nuclear Energy Cost Assumptions 34
Industrial Processing to Reduce CO_2 Emissions 36
31 Nuclear-Integrated Processing Platforms 49
Material Handling Costs 58
Large Material Handling Equipment 66
Independent Order-of-Magnitude Check on Program Scale 75
Standards Development 84
Direct Air Capture (DAC) Using Nuclear Power 96
Automated Construction of Thermal Storage Tanks 100
Electrified Rail to Assist in Site Construction 110
Different Ways to Make Ammonia 122
Ships Powered by Ammonia 128
Ammonia to Electricity Conversion Efficiency (e.g. on a ship) 136
Cost of Ammonia Fuel Cell on Large Ships 141
Liquid Ammonia Production Quantity 147
Purchase Orders 152
Purchase Order Recipients 156
Project Management 158

Funding Source Requirements 160
Top Industrial Processing and Nuclear Integration References 163
Top Researchers and Organizations 170
Scientific and Engineering Advisory Panel..... 172
Management Advisory Panel..... 174
Agreements..... 175
Key Codes, Standards, and Data Sources 177

Nuclear-Integrated Industrial Processing Platforms

In the proposed concept, large platforms of industrial processing equipment are co-located with nuclear power plants and powered by heat transported through isolated steam, hot-water, or other qualified intermediate loops. The reference transport envelope is 12 m wide, with 24 m modular increments: one to four modules can form 12 × 24, 12 × 48, 12 × 72, or 12 × 96 m platforms. The largest four-module platforms can be transported by double rail and positioned side-by-side near the reactor (Figure 1). Site-wide efficiency can be improved by cascading useful heat among processes, recovering waste heat, and using thermal storage, heat pumps, or mechanical-vapor recompression where economical. The platform families include generic industrial heat services, direct air capture, hydrogen and ammonia production, and sector-specific industrial processing.

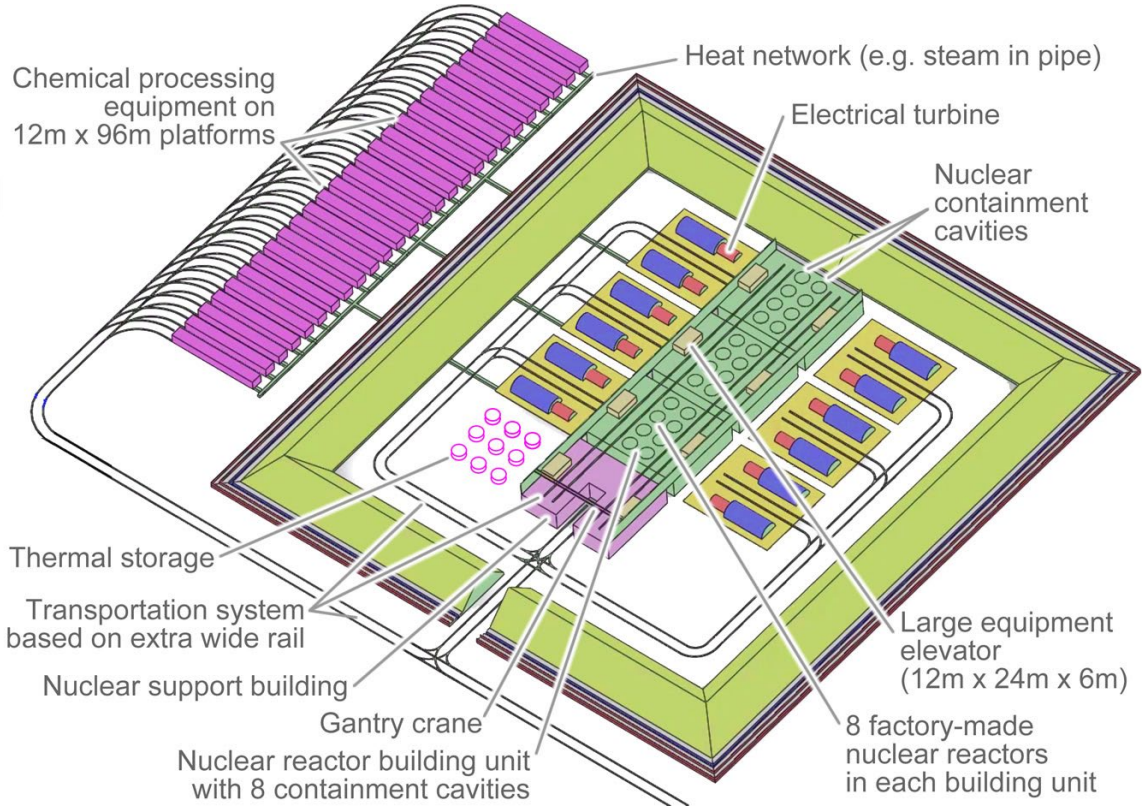

Figure 1: Nuclear reactors co-located with multiple platforms of chemical processing equipment. Concept illustration by G. Weinreb.

Industrial Processing Yard

In theory, chemical processing platforms could be positioned side-by-side, as shown below. In this illustration, platforms are shown in green, pipes carrying heat from the nuclear reactors are shown in blue, underground tunnels that support pipes in between platforms are shown in light violet, and

underground tunnels that support moving material into and out of the chemical processing platforms are shown in light purple.

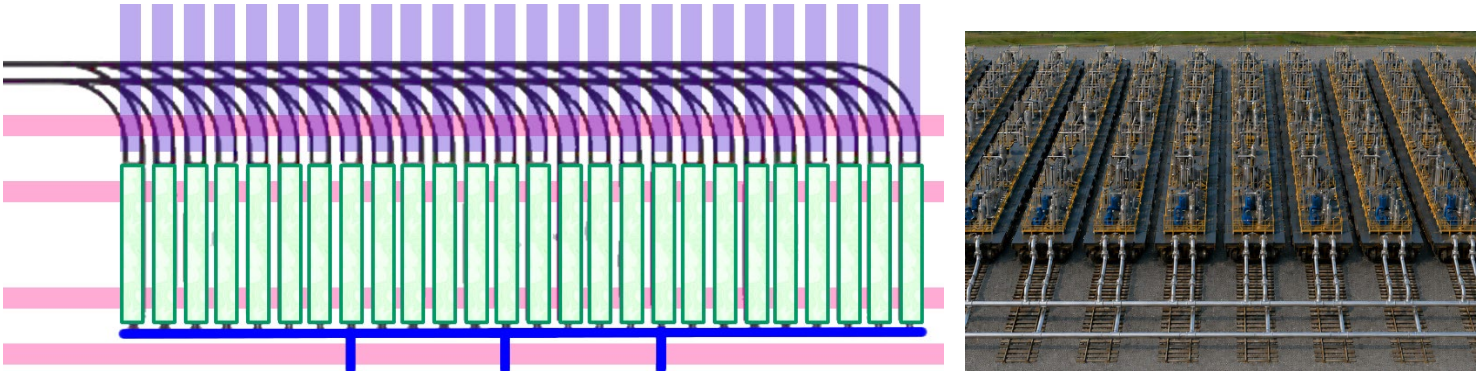

Figure 2: Chemical processing yard consists of processing platforms side-by-side.

There are two primary approaches to underground tunnel construction. One method involves manufacturing precast concrete modules at an off-site factory, transporting them to the site, and joining them with more concrete. Alternatively, rebar cages and forms could be assembled off-site, transported to the site, and then filled with wet concrete. Specialized equipment and factories could be developed that support both methods.

Platform Transporter

In theory, equipment could be developed that moves 12m × 96m platforms across land via double-rail, and across water via specially built ships. To prevent bending when traveling over hills or around curves, these platforms could be mounted on a ridged truss supported by jacks between the truss and railcars, as illustrated below.

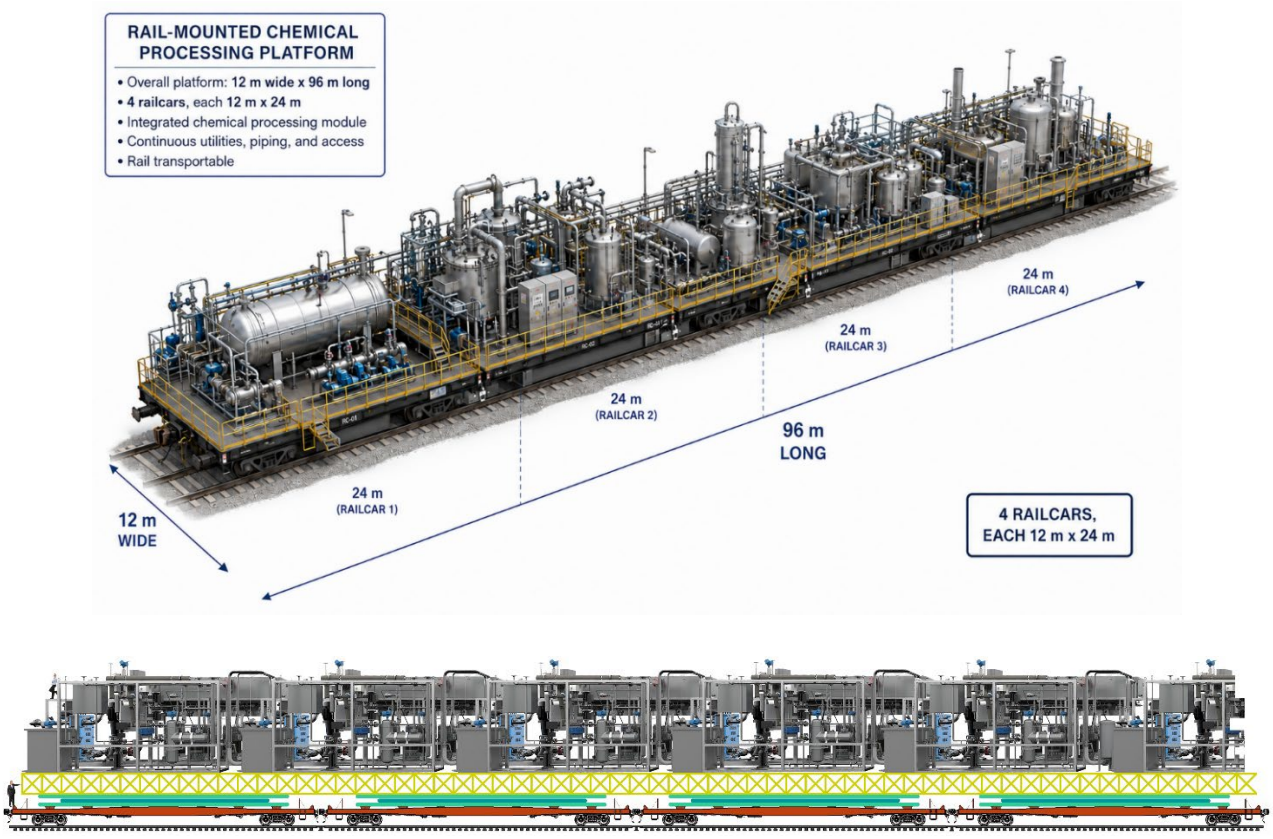

Figure 3: Concept illustration of chemical processing platform transported on a rigid truss, mounted on four double-width railcars.

Thermal Storage

When electricity demand is low, nuclear reactors can increase the temperature of a hot liquid in a thermal storage tank; and when electricity demand is high, the stored energy can be used to make electricity. In other words, to minimize costs, thermal storage can bridge the gap between fluctuating demand and constant supply. Also, thermal storage can minimize waste heat by consuming it when available at a temperature higher than the storage temperature. In theory, R&D could reduce thermal storage costs by supporting the development of equipment that automates the construction of thermal storage tanks.

China Strategy

Cost and engineering discipline. China is the controlled industrial reference because the program can compare standardized nuclear heat/electricity and modular industrial equipment against large, operating fossil-based industrial supply chains rather than against an abstract global average. The cost

model should not apply a generic “China discount.” For each platform family it should reconstruct energy use, throughput, equipment CAPEX, factory fabrication, installation, maintenance, consumables, labor, logistics, financing, product specification, emissions, and utilization using the same boundaries for the nuclear and fossil cases.

Scale and global applicability. The industrialization model should show how successful pilot platforms become repeated products: qualified shipyards and factories, module transport and port capacity, standardized bays and utilities, spare-parts systems, calibration and conformance laboratories, software support, and trained operators. The $1 billion R&D program develops and validates the platform families and enabling systems; subsequent industrial capital finances commercial factories, nuclear-industrial sites, and ship fleets. Proven platform designs, interface standards, process specifications, software, transport equipment, thermal systems, and QA/test methods can then be adapted in Europe, the United States, and elsewhere, subject to local reactor availability, industrial codes, labor, environmental and process-safety requirements, transport geometry, and market conditions.

Relationship to the Nuclear Power Automated Site-Construction Initiative

The $1B automated site-construction initiative and the $1B chemical processing initiative (this document) should be treated as adjacent layers of one industrial system, not as two programs that independently buy the same cranes, transporters, rail systems, digital tools, and test facilities. The nuclear site-construction initiative develops the common construction backbone needed to build HTR-PM and Hualong One/PWR plants. The industrial-processing initiative develops the process-yard hardware that uses the resulting nuclear heat and electricity.

A simple rule should govern procurement: if the site-construction initiative has already qualified a shared asset to the load, precision, safety, and interface requirements needed by the industrial yard, the industrial-processing initiative should use that asset rather than fund a second version. The industrial-processing initiative pays only for incremental industrial requirements—for example a 12 × 96 m processing-platform payload, ship/shore transfer, process-material connections, or additional operating cycles that are not part of the reactor-construction mission.

Division of scope

Capability	Site-construction initiative	Industrial-processing initiative	Coordination rule
Heavy-module transport	Develop common factory/port/site heavy-transport architecture for reactor and structural modules.	Qualify incremental 12 × 96 m industrial-platform payloads, rigid-truss behavior, repeated platform replacement, and process-yard routes.	One transporter family where practical; the industrial-processing initiative pays only for added payload/interface requirements.
Heavy lift, gantry, crane and positioning	Develop common heavy-handling machines, controls, metrology, proof-load methods, and reactor-specific fixtures.	Develop industrial-platform lifting/docking fixtures and bay interfaces.	Reuse the common gantry/crane platform; do not buy duplicate base machines.
Rail and site logistics	Develop construction rail/site-logistics geometry and heavy movement interfaces used by HTR-PM/PWR construction.	Develop segmented electrified-rail power delivery, collectors, switching, lockout, warning systems, and industrial-yard operating profiles.	Share rail geometry and safety envelope; the industrial-processing initiative owns the electrification extension unless the site-construction initiative adopts it first.
Civil works, tunnels and material flow	Develop reactor-site excavation, concrete, rebar, structural automation, spoil removal, and construction logistics.	Develop pre-equipped industrial tunnels, material-distribution spines, manifolds, conveyors, storage and waterfront modules.	Share civil automation methods; the industrial-processing initiative pays for process-yard-specific tunnel contents and logistics.
Digital twin, autonomy and	Develop regulated construction digital twin, machine control,	Extend the same digital architecture to industrial modules, interoperability,	Use common identifiers/data protocols where possible; avoid

metrology	metrology, configuration and QA data thread.	maintenance, process interfaces and conformance testing.	parallel software stacks.
Nuclear/industrial heat boundary	Define plant-side construction envelopes and reactor-side interfaces needed for the reference plants.	Develop qualified isolation loops, heat grades, steam/molten-salt interfaces, metering, process-yard headers and platform connection standards.	The industrial-processing initiative owns the non-nuclear process interface; the reactor safety boundary remains controlled by the nuclear design authority.
Thermal storage and industrial utilities	May provide site construction methods and shared heavy equipment if useful.	Develop automated tank-construction machines, storage modules, heat pumps/MVR, utility platforms, and industrial distribution systems.	Use site-construction initiative machines when qualified; fund only incremental industrial tooling.
Industrial processes, DAC and ammonia	Outside the site-construction initiative except for generic site interfaces.	Own the 31 platform families, DAC, hydrogen/ammonia, marine ammonia fuel cells, process pilots and long-duration industrial testing.	No overlap: these are industrial-processing deliverables.

Funding and interface rule

- Fund each shared physical asset only once; charge the second program only for incremental requirements, tooling, tests, or operating time.
- Maintain one interface-control document for module dimensions, rail geometry, lifting points, electrical power, heat, process piping, controls/data, metrology, and safety boundaries.
- Use a joint configuration-control board for changes that affect both reactor construction and the industrial yard.
- Maintain a procurement ledger marking every major package as construction-only, industrial-only, or shared, with the funding split and acceptance tests stated before award.

This division lets the two $1B programs reinforce each other. The site-construction initiative proves that standardized Chinese PWR and HTR-PM plants can be built with a more automated site system; the industrial-processing initiative then turns low-cost nuclear heat and electricity into a standardized industrial campus without paying a second time for the same construction backbone.

Global Deployment Economics

Purpose and Scope

The controlled technical and economic reference remains the China nuclear-industrial system: Automated China PWR for low-cost electricity and low/medium-temperature heat, HTR-PM for applications that genuinely need higher-temperature heat, and the 31 standardized platform families plus their shared logistics and interface architecture. This chapter adds secondary United States and European deployment sensitivities. Their purpose is not to create separate $1 billion hardware programs or to assume that Hualong One, HTR-PM, or a Chinese chemical-plant approval can simply be transplanted. Instead, the same measured platform evidence should be re-costed and re-qualified against locally licensable nuclear heat sources, industrial codes, fossil alternatives, supply chains, and market conditions.

China therefore remains the primary proving environment: it asks whether repeated factory-built platforms can reduce delivered cost and CO2 emissions under a high-volume industrial reference case. The regional cases ask a second question: after the platform technology and standards mature, how much of that advantage survives when local energy prices, labor, financing, codes, environmental rules, transport constraints, tariffs, feedstocks, and existing industrial assets are included?

Regional Economic Comparisons

China. Use the existing China chapters as the controlled baseline. Compare each selected platform with both relevant new fossil construction and the going-forward or marginal cost of existing fossil industrial capacity, using the same product specification, throughput, energy boundary, logistics boundary, and financing convention. The Automated PWR and HTR-PM energy values remain planning scenarios until the companion construction program and later operating evidence support them.

United States. Map the validated industrial platforms to locally licensable nuclear sources rather than assume the Chinese reactor designs are directly deployable. For low- and medium-temperature process heat, compare nuclear heat, nuclear-electric heat pumps/MVR, and other platform routes primarily with natural-gas boilers, furnaces, and existing gas-intensive industrial assets, while retaining coal or oil where they are actually the incumbent. Re-price U.S. engineering and construction labor, module fabrication, utilities, port and heavy transport, process-safety compliance, environmental permitting, nuclear-interface work, financing, insurance, feedstocks, product delivery, and maintenance. Hydrogen, ammonia, DAC, and marine-fuel cases should be compared on delivered product or service cost—not solely on reactor LCOE.

Europe. Use representative-country and industrial-cluster cases rather than one permanent “Europe” cost number. Reactor availability, gas and electricity prices, industrial heat sources, financing, port

access, carbon costs, and policy support vary substantially across Europe. Re-price the same platform evidence for selected countries or clusters and separate the underlying engineering economics from EU ETS costs, grants, contracts-for-difference, state financing, hydrogen support, or other policy instruments. For high-temperature HTR-like applications, treat suitable locally licensable reactor heat as a conditional input rather than assume that 300–500°C nuclear heat is immediately available everywhere.

Platform and Supply-Chain Localization Scenarios

Path 1: China-manufactured platform systems. Mature process platforms, utility modules, selected transport equipment, controls, and standardized components are manufactured by qualified Chinese suppliers and exported where size, route, tariff, security, code, and local-content constraints permit. The destination country performs local foundations, nuclear/industrial integration, operator qualification, maintenance, inspections, and jurisdiction-specific acceptance.

Path 2: Transferred designs with local production. Sponsor-funded platform designs, manufacturing packages, software, interface specifications, conformance tools, and test procedures are replicated by U.S. or European suppliers to the maximum extent allowed by intellectual-property, export-control, cybersecurity, and nuclear-information rules. Local shipyards and heavy industrial fabricators may be especially important where a 12 × 96 m module cannot economically cross inland transport networks.

Path 3: Hybrid localization. High-value standardized modules and components are imported where economical and lawful, while very large structures, safety-sensitive equipment, locally competitive machinery, local-content items, and site-specific interfaces are manufactured domestically. The mix should be selected by delivered lifecycle product cost and qualification risk rather than by an assumed localization percentage.

For all three paths, regional economics should include platform CAPEX and amortization, shipping, port handling, tariffs, local-content rules, installation, spare parts, maintenance, software/cybersecurity adaptation, operator training, pressure-equipment and process qualification, nuclear-interface work, insurance, financing, utilization across a multi-project pipeline, and the cost of adapting platform connections to the local reactor and industrial site.

Nuclear, Industrial, and Standards Qualification Transfer

Performance proven at a non-nuclear China pilot does not automatically establish permission to connect the same platform to a licensed U.S. or European nuclear site or to operate an ammonia, hydrogen, high-temperature, marine, or hazardous-chemical system under another jurisdiction. The transfer package should separate three layers: (1) generic platform capability and open interfaces that can be reused broadly; (2) reactor- and site-specific heat/electricity connections and safety boundaries that must be re-engineered; and (3) jurisdiction-specific nuclear, pressure-vessel/piping,

electrical, machinery, hazardous-area, environmental, worker-safety, cyber, marine/classification, and permitting requirements. The regional economic model should charge adaptation and qualification cost rather than treat it as free.

The standards program should therefore develop a common international interface core plus jurisdictional profiles. U.S. implementation can map the platform core to applicable ASME, API, NFPA, IEEE/ISA/IEC and other requirements; European implementation can map it to applicable EN/CEN/CENELEC/IEC standards and EU or national pressure-equipment, hazardous-area, machinery, environmental, and nuclear requirements. The objective is not identical regulation, but enough interface stability that suppliers can reuse equipment without redesigning the entire industrial site.

Regional Engineering Support

Major full-scale prototype and endurance work can remain concentrated in the primary China-centered program while smaller U.S. and European teams participate from the beginning in standards mapping, locally licensable reactor/heat-source interfaces, process-safety reviews, industrial cost models, port and transport studies, and selected interoperability tests. This avoids duplicating expensive hardware while reducing the risk that a China-optimized platform is discovered in Year 5 to be difficult to qualify or manufacture elsewhere. Area 1 and Area 9 should explicitly fund this regional engineering participation within their existing budgets.

Public R&D, Carbon, Subsidy, and Policy Treatment

Each regional comparison should show two economic views. The engineering-economics view reports the full platform, energy, logistics, financing, and product cost without hiding carbon prices, tax credits, grants, direct plant subsidies, concessionary financing, or similar policy support inside the technology cost. The policy-adjusted view then adds the actual carbon costs, tax incentives, grants, tariffs, local-content rules, hydrogen or industrial-decarbonization support, and other policies applicable to the project. Sponsor-funded R&D may be treated as sunk cost in a project-owner deployment view, but the fully loaded societal view should disclose those public contributions separately.

Program Deliverable

Area 1 should maintain a Global Deployment Economics and Transferability workbook linked to the controlled China models. By Year 5 it should report, for China and selected U.S. and European cases: platform/product; locally licensable nuclear or heat-source basis; fossil comparator; energy and feedstock prices; localization path; factory/site CAPEX; labor and utilization; logistics; maintenance; qualification/adaptation; financing; delivered product cost; CO2 reduction; policy adjustments; and uncertainty. China remains the controlled baseline; non-China cases are transparent sensitivities using the same measured platform evidence.

Proposed $1B R&D Initiative

Prompt to AI: Summarize a 5-year $1B R&D program to co-locate large, factory-built chemical processing platforms with nuclear power plants. This is based primarily on the "Co-locate Chemical Processing with Nuclear Power" chapter within the Climate Moonshot Business Plan.

The program should develop ≥ 30 standardized industrial-processing platforms that support ammonia production, along with other common chemical processing tasks. Also, it should: (a) develop marine ammonia fuel-cell systems, develop common standards, develop electrified construction rail standards, and develop machines that automate thermal-storage-tank construction. Prototypes should be built and tested for long-durations without nuclear fuel. Focus on PWR and HTR-PM nuclear power plants built in China with automated site construction.

Response From AI:

The five-year R&D program is centered on a simple idea: manufacture large industrial-processing equipment in factories or shipyards, move it to a prepared nuclear industrial site as standardized modules, connect it through open mechanical, thermal, electrical, material-handling, safety, and data interfaces, and prove the system with progressively larger prototypes before committing to commercial deployment.

The program uses two Chinese nuclear reference technologies. Automated China PWR is the preferred source for very-low-cost electricity and low/medium-temperature heat below roughly 300°C; while Automated China HTR-PM is the preferred source when 300–500°C-class direct process heat materially improves the process. The plant architecture should cascade heat from higher-temperature platforms into lower-temperature platforms, with industrial heat pumps and mechanical-vapor recompression recovering waste heat whenever practical.

A portfolio of 31 standardized industrial-processing platform families is developed using the 12 m-wide, 24 m-increment module standard. The first five generic thermal platforms receive disproportionate attention: industrial heat-pump steam generation, mechanical-vapor-recompression evaporation, direct PWR process steam, distillation/reboiler service, and HTR-PM 300–500°C process heat. Sector-specific platforms then adapt those building blocks to pulp and paper, food, textiles, alumina, chemicals, refining, metals, glass, ceramics, minerals, and other large fossil-fired heat loads.

The fund also retains dedicated hydrogen and ammonia development because ammonia can serve fertilizer markets and, potentially, marine propulsion. It develops factory-buildable HTSE / SOEC hydrogen and Haber-Bosch ammonia platforms, then carries three shipboard high-temperature fuel-cell architectures through R&D: direct/internal-cracking SOFC, partial-external-cracking SOFC, and externally cracked PCFC. Small systems are run for thousands of hours with teardown inspections so

stack degradation, catalyst changes, ammonia slip, seals, piping corrosion, thermal cycling, and balance-of-plant aging are measured rather than assumed.

The supporting construction system is treated as part of the product. The program develops giant module transport, modular waterfront equipment, tunnels, manifolds, conveyors, storage, a common material-distribution spine, segmented electrified rail for construction machinery, and specialized machines that automate the construction of very large thermal-storage tanks. Standards are developed early so equipment from different suppliers fits together physically and communicates reliably.

Most R&D testing is non-nuclear. Fossil-fired heaters provide PWR-like and HTR-PM-like heat to qualified intermediate loops. Site choreography may be tested at 1:100 scale, transport systems around 1:10, process platforms at pilot throughput rather than simple geometric scale, and pipes/valves/corrosion loops at representative diameter, chemistry, pressure, and temperature. This preserves the real degradation mechanisms that a tiny geometric model would miss.

Key boundaries of the fund

- **Included:** FEED/reference designs for all 31 platform families; substantial integrated pilots for roughly the highest-priority 8–12 families; subsystem rigs or representative test articles for the remainder; ammonia platforms; shipboard ammonia fuel-cell prototypes; giant-module logistics; material handling; thermal storage; electrified construction rail; open standards; interoperability testing; full-temperature flow/corrosion loops; digital engineering; cost models; safety analysis; and long-duration non-nuclear operation.
- **Excluded:** construction or operation of a licensed commercial nuclear plant, nuclear fuel work, commercial-scale fertilizer production, construction of a commercial ship fleet, or a guarantee of any specific electricity, heat, ammonia, or shipping cost.
- **Reference energy economics:** use the Automated China PWR and HTR-PM planning ranges in the enclosed nuclear site-construction initiative, and clearly label them as scenario/R&D target estimates rather than demonstrated commercial costs.
- **Safety principle:** preserve a qualified isolation boundary between nuclear systems and ordinary equipment; design industrial yard and marine systems so they do not compromise safety.

Work package	Budget	Share	Principal five-year deliverable
Area 1. Reference architecture, 31-platform portfolio	$70M	7.0%	Freeze the dual-reactor industrial-site architecture, 31 platform reference designs, heat-cascade logic, interfaces, hazard boundaries, digital model, commercial screening methods, and China/U.S./European deployment economics and transferability sensitivities.
Area 2. 31 industrial-processing platform designs and pilot prototypes	$180M	18.0%	Design all 31 platform families; build substantial integrated pilots for roughly 8–12 highest-priority families; concentrate endurance testing on the first five generic thermal platforms, the DAC train, and other high-risk systems. Ammonia prototype is primarily funded under Work Package 3.
Area 3. Hydrogen and ammonia production platforms	$95M	9.5%	HTSE/SOEC hydrogen, air separation/nitrogen, Haber-Bosch ammonia, refrigeration/storage/loading, heat recovery, and a repeatable factory-buildable ammonia platform package.
Area 4. Marine ammonia fuel-cell propulsion prototypes and endurance testing	$70M	7.0%	Competing direct-SOFC, partial-cracking SOFC, and cracked-ammonia PCFC systems; small integrated propulsion prototypes; thousands-of-hours durability testing and teardown inspection.
Area 5. Material handling, tunnels, storage, waterfront, and modular logistics equipment	$85M	8.5%	Pre-equipped tunnels, pipe/manifold modules, pumps/compressors, conveyors, transfer towers, storage modules, waterfront loading/unloading, and the common distribution spine.
Area 6. Heavy module transport, double rail, ship/shore transfer, and gantry systems	$65M	6.5%	Rigid truss, synchronized carriers, 12×24 through 12×96 m module transport, temporary track, loading, ship transfer, alignment, recovery, and repeated handling demonstrations.
Area 7. Electrified rail and high-power construction-machine interface	$40M	4.0%	Segmented LV/MV enclosed conductor rail, collectors, switchgear, grounding, interlocks, regenerative operation, warning/lockout systems, and a full-voltage reduced-length test loop.
Area 8. Automated construction of thermal-storage tanks	$50M	5.0%	Tank-jacking, shell-panel placement, automated welding/NDE, floor/roof/insulation handling, modular pump/manifold equipment, and representative tank construction demonstrations.
Area 9. Standards, interoperability lab, digital thread, and conformance testing	$55M	5.5%	Open mechanical, dimensional, heat, piping, electrical, rail, controls/data, cybersecurity, waterfront, storage, and digital-thread standards plus multi-vendor interoperability tests.
Area 10. Integrated prototype campus, flow/corrosion loops, and long-duration campaigns	$65M	6.5%	Non-nuclear test campus with PWR-like and HTR-PM-like heat loops, full-temperature materials/corrosion systems, site choreography, integrated platform campaigns, maintenance, and reliability growth.
Area 11. Program management, safety, techno-economics, and independent review	$25M	2.5%	Owner-side systems engineering, financial control, HAZOP/functional safety, cost and emissions models, independent verification, public data release, and Phase-II decisions.
Area 12. Central contingency and risk reserve	$200M	20.0%	Released only for documented technical, supplier, standards, safety, test, regulatory, or integration risks after independent review.

Program Architecture and $1 Billion Budget

Prompt to AI: Allocate a fixed $1 billion over five years to implement the accepted discussion-2b recommendations. The fund must cover ≥ 30 platform designs, prototypes, standards, ammonia production, ammonia-to-electricity ship systems, long-duration testing, material handling, electrified rail, automated thermal-storage-tank construction, and integrated demonstrations while retaining a centrally controlled risk reserve.

Response From AI:

The recommended program is $800 million of planned technical work plus a $200 million centrally controlled reserve. The reserve is deliberately large because pressure systems, high-temperature electrochemistry, large moving structures, marine operations, tank welding, electrified rail, and long-duration durability testing can expose expensive integration failures. The technical program is not divided equally among 31 platforms; generic thermal platforms and shared infrastructure receive more money because they can serve many industries.

Suggested five-year funding profile

Year	Technical work	Reserve	Total	Primary emphasis
Year 1	$120M	$20M	$140M	Architecture; 31-platform FEED; standards v0.9; fuel-cell and ammonia concept competitions; test-campus and rail/tank requirements.
Year 2	$175M	$30M	$205M	Preliminary design; first platform pilots; standards v1.0; corrosion loops; fuel-cell stacks; logistics and electrified-rail component tests.
Year 3	$220M	$40M	$260M	Peak detailed design/fabrication; ammonia pilot; transport/logistics prototypes; rail loop; tank-building machines; integrated test-site commissioning.
Year 4	$190M	$60M	$250M	Long-duration campaigns; repeated platform and logistics operation; thermal-storage tank build; multi-vendor interoperability; ship propulsion endurance runs.
Year 5	$95M	$50M	$145M	Reliability growth, teardowns, replicated cost models, standards v2.0/international packages, China deployment designs, U.S./European transferability cases, and Phase-II readiness.
Total	$800M	$200M	$1.000B	Reserve remains sponsor-controlled and may remain unspent.

Program-level success criteria

- Thirty-one controlled industrial-processing platform-family reference designs using the 12 m-wide modular envelope, including a clear priority order and common utility/heat interfaces.

- FEED/reference designs for all 31 platform families; substantial integrated pilots for roughly the highest-priority 8–12; subsystem rigs or representative test articles for additional viable families where needed; and no requirement to prototype weak cases that fail the economic, climate, or technical screen.
- An integrated hydrogen-to-ammonia pilot and manufacturing package using both PWR- and HTR-PM-derived heat/electricity assumptions where appropriate.
- Three competing marine ammonia fuel-cell propulsion architectures tested with real ammonia, electric loads/motors, heat recovery, emissions control, and teardown-based degradation measurement.
- Demonstrated 12×24 through 12×96 m module handling, ship/shore transfer, modular tunnels/material systems, electrified-rail machine power, and automated thermal-storage-tank construction.
- Open interface standards stable enough that at least three independent supplier teams can connect modules at the interoperability laboratory without bespoke redesign.
- Independent, reproducible cost models that use the enclosed Automated China PWR/HTR-PM LCOE assumptions as explicit scenarios and compare both new-build and existing-fossil marginal-cost cases. Include selected U.S. and European transferability cases using the same measured platform evidence, with local energy/feedstock prices, localization, qualification, financing, policy, and product-delivery boundaries stated explicitly.

What the fund would not promise

- It would not promise that all 31 platform designs become commercial products; the program is designed to down-select and concentrate later spending on systems that prove useful and economical.
- It would not build 31 near-commercial-scale demonstration plants. The program differentiates FEED/reference designs, integrated pilots, subsystem rigs, and long-duration materials/component tests according to the physics and risk of each platform family.
- It would not guarantee $365–375/t ammonia, $31/MWh PWR electricity, $37/MWh HTR-PM electricity, or any other scenario estimate. Those values are inputs to be tested, not program outcomes.
- It would not claim a 1:100 model can establish corrosion, heat transfer, pressure-drop, fuel-cell life, catalyst life, or industrial reliability. Those require representative physical conditions and endurance testing.
- It would not connect prototypes to nuclear fuel during this five-year R&D program unless a separately licensed host-site project later chooses to do so.

Program Crosswalk

Prompt to AI: Show how the work packages map to platform families, enabling systems, principal deliverables, five-year budgets, and the stage-gate architecture.

Response From AI:

This crosswalk connects the controlled $1 billion budget to the technical portfolio and decision process. Platform, prototype, standards, logistics, ammonia, and marine analyses later in the proposal are explanations of these work packages rather than additional budget lines. The same scope must not be counted twice.

Work package	Budget	Principal five-year deliverable	Platform / system connection	Gate focus
Area 1. Reference architecture / digital twin	$70M	Dual-reactor site architecture, 31 reference designs, heat cascade, interfaces, hazards, digital model, economic screening, and U.S./European transferability sensitivities.	All 31 + sitewide utility hub	G0–G2
Area 2. Industrial platform designs / pilots	$180M	FEED for all 31; substantial pilots for roughly 8–12 priority families; targeted rigs/test articles for additional viable cases.	Platforms 1–31	G1–G5
Area 3. Hydrogen and ammonia production	$95M	HTSE/SOEC, nitrogen, Haber-Bosch, storage/loading, heat recovery, and factory-buildable ammonia package.	Platforms 28–31	G1–G5
Area 4. Marine ammonia fuel cells	$70M	Competing SOFC/PCFC architectures, integrated propulsion prototypes, long-duration runs, emissions control, and teardown evidence.	Downstream ammonia use	G1–G5
Area 5. Material handling / waterfront	$85M	Tunnels, manifolds, pumps/compressors, conveyors, storage modules, waterfront equipment, and distribution spine.	All site process families	G1–G5
Area 6. Heavy module transport	$65M	Rigid truss, synchronized carriers, 12 × 24 through 12 × 96 m movement, ship/shore transfer, alignment, and recovery.	All modular platforms	G1–G5
Area 7. Electrified rail	$40M	Segmented LV/MV conductor rail, collectors, switchgear, grounding, interlocks, warning/lockout, and high-power tests.	Construction / handling machines	G1–G5
Area 8. Thermal-storage construction	$50M	Tank jacking, automated panel handling/welding/NDE, insulation, pumps/manifolds, and representative tank build.	Sitewide thermal storage	G1–G5

Area 9. Standards / interoperability	$55M	Open interfaces, data/cyber schemas, conformance tools, and multi-vendor plug tests.	All modules and shared systems	G0–G5
Area 10. Integrated prototype campus	$65M	PWR-like/HTR-like heat loops, corrosion rigs, integrated logistics/platform campaigns, maintenance, and reliability growth.	Selected priority platforms	G3–G5
Area 11. Program management / independent review	$25M	Systems engineering, HAZOP, techno-economics, lifecycle emissions, audit, replication, and final Phase II decisions.	Program-wide	All gates
Central risk reserve	$200M	Sponsor-controlled response to documented technical, supplier, standards, safety, test, regulatory, or integration risk.	Not pre-awarded	Written gate/change approval

Stage Gates and Go/No-Go Criteria

Prompt to AI: What explicit stage gates should control progression from the industrial-platform architecture through down-selection, prototype fabrication, integrated endurance testing, interoperability, and any later Phase II scale-up?

Response From AI:

Passing a gate authorizes only the next R&D increment. It does not require that a platform survive the portfolio, and it does not authorize a licensed nuclear-industrial complex or commercial ship fleet. Gate evidence should be fixed before major commitments and reviewed independently.

Gate	Required evidence	Decision enabled	If not met
G0. Mission / economic baseline (M6–9)	Controlled PWR/HTR energy scenarios; fossil new-build and existing-asset comparators; product boundaries; CO_2 metric; module standard; initial hazard and market screen.	Authorize detailed architecture and competing concepts.	Revise assumptions or remove applications that lack a credible climate/economic case.
G1. Architecture / interfaces ready (M12)	31 FEED concepts; utility/heat classes; 12 m × 24 m-increment geometry; isolation boundaries; standards v0.9; test methods; design ownership.	Authorize bench/subscale risk reduction and supplier competitions.	Close interface/data gaps; do not fabricate against unstable requirements.
G2. Platform down-selection (M18–24)	Comparable economics, tCO_2 avoided, market scale, technical risk, modularity, maintainability, and early test evidence for all platform families.	Select roughly 8–12 major integrated pilots; retain targeted rigs for other viable cases.	Stop at FEED, redesign, or defer weak cases rather than spreading hardware money evenly.
G3. Major prototype authorization (M30–36)	Controlled design; HAZOP/functional safety; suppliers; test-site interfaces; instrumentation; acceptance criteria; cost/schedule update; independent review.	Commit major prototype and campus fabrication.	Delay or narrow fabrication; correct safety, interface, or evidence deficiencies first.
G4. Integrated / endurance readiness (M40–46)	Qualified hardware; calibrated measurements; operating envelope; maintenance/recovery plan; materials/corrosion baseline; open deficiencies reviewed.	Begin long-duration integrated operation and repeated module/logistics campaigns.	Correct deficiencies; do not waive material safety or measurement requirements to preserve schedule.
G5. Replication / interoperability (M50–57)	Repeated performance, degradation/teardown evidence, multi-vendor conformance, independent cost reconstruction, and reproducible data package.	Freeze mature designs/standards and prepare deployment-grade engineering packages.	Treat results as unconfirmed; redesign, repeat, or stop the workstream.

G6. Phase II decision (M60)	Platform-by-platform technical, economic, climate, reliability, safety, and market evidence with unresolved risks explicitly stated.	Recommend proceed / redesign / stop for each platform family and enabling system.	No Phase II entitlement; redirect remaining work to higher-value evidence or close out.

Program-level stop rule

Stop or redirect a workstream when repeated tests cannot establish a credible delivered-cost or CO_2 advantage, when safety or environmental requirements cannot be met, when durability or maintainability is inadequate, when an interface cannot be standardized economically, or when the expected value of additional information is lower than competing uses of the remaining funds.

Program-Wide Measurement and Acceptance Criteria

Prompt to AI: What common measurements and acceptance evidence should every major platform or enabling-system workstream provide so the program can compare technical performance, economics, climate benefit, reliability, safety, and modularity on the same basis?

Response From AI:

Every major prototype should be judged by a common evidence spine. A system that works technically but cannot beat the relevant fossil counterfactual, cannot survive long-duration operation, or cannot be reproduced independently should not advance merely because fabrication is complete.

Metric	Required measurement	Acceptance evidence	Used for
Energy / process performance	MWh_e and MWh_{th} per tonne; COP/MVR ratio; heat recovery; throughput; product quality; turndown.	Calibrated mass/energy balances under controlled operating conditions with uncertainty.	G2–G6 platform comparison
Economics	Platform CAPEX, installation, O&M, maintenance, consumables, financing, utilization, and delivered product/heat cost.	Independent cost reconstruction against both new fossil construction and existing-fossil marginal operation.	Down-selection and Phase II
Climate benefit	Direct and lifecycle tCO_2 avoided or removed per tonne product; $/$tCO_2$; system boundary and counterfactual.	Reproducible lifecycle inventory plus measured energy/material inputs; no double counting across platforms.	Portfolio ranking
Reliability / durability	Availability, forced outages, degradation, corrosion, catalyst/stack life, leak rate, cycling, maintenance interval.	Long-duration runs, inspections/teardowns, failure log, corrective actions, and reliability-growth record.	G4–G6
Safety / environment	HAZOP findings, hazardous-area controls, NH_3/H_2 release, electrical/rail hazards, marine risks, emissions, waste, water.	Independent closure or explicit acceptance of critical findings; functioning shutdown/containment and incident reporting.	Every major test gate
Modularity / logistics	Factory-completion fraction, transport limits, installation/removal time, docking/alignment, reconnect time, recovery from faults.	Repeated measured module moves and replacement campaigns using controlled geometry and load cases.	G3–G6
Interoperability	Mechanical, heat, electrical, controls/data, cyber, rail, material, and utility interface conformance.	At least three independent supplier teams connect at the interoperability lab without bespoke site redesign.	G5 standards freeze
Independent reproducibility	Raw data, calibration, source files, test configuration, exclusions, assumptions, failures, and uncertainty.	Unaffiliated team can reconstruct key performance, cost, and climate claims from the delivered record.	Public claims and G6

Portfolio down-selection rule

The program should deliberately spend unevenly. All 31 platform families receive enough engineering to make a comparable decision, but major hardware funding is concentrated only after the G2 portfolio review.

Portfolio tier	Minimum deliverable	Spending rule
All 31 platform families	FEED/reference design, interfaces, safety screen, cost/CO_2 model, and priority score.	Required before G2; hardware is not automatic.
Highest-priority roughly 8–12	Substantial integrated pilots plus representative endurance/maintenance testing.	Receive most WP2 prototype spending after G2.
Other viable families	Subsystem rig or representative test article only where a specific uncertainty justifies it.	Fund to resolve the decision question, not to create equal prototypes.
Weak / uneconomic cases	Documented negative result and reusable design/economic record.	Stop at FEED or after limited risk reduction; redirect funds.

Development Phases

Prompt to AI: For each work package in the $1B fund, identify the development phases, approximate cost allocation, deliverables, and the major decision gates. Ensure prototypes, standards, corrosion/endurance testing, ammonia fuel cells, electrified rail, and automated thermal-storage-tank construction are explicit.

Response From AI:

Common stage	Timing	Purpose
A. Requirements and feasibility	Months 0–12	Freeze mission, interfaces, hazards, economics, reference loads/temperatures, test methods, and vendor competition.
B. Preliminary design and risk reduction	Months 6–24	Down-select concepts; build subscale/bench rigs; release standards v0.9/v1.0; close highest-risk physics and safety questions.
C. Detailed design and procurement	Months 18–36	Controlled drawings, software, fabrication packages, supplier qualification, test procedures, and long-lead procurement.
D. Prototype fabrication and qualification	Months 24–46	Build and qualify process, transport, rail, tank, ammonia, fuel-cell, and infrastructure prototypes.
E. Integrated and endurance demonstrations	Months 36–57	Operate systems together; months/years-equivalent cycling where practical; inspect degradation; repeat failed tests after redesign.
F. Reliability growth and handoff	Months 48–60	Freeze replicable designs and standards, quantify remaining risks, and prepare first-site or first-ship Phase-II packages.

These A–F phases describe execution inside each work package. Program spending authority is governed by the standalone G0–G6 stage gates above; the work-package hold points below add only technology-specific safety, readiness, or performance conditions.

Area 1. Reference architecture, 31-platform portfolio, and digital twin - $70M

Freeze the dual-reactor industrial-site architecture, 31 platform reference designs, heat-cascade logic, interfaces, hazard boundaries, digital model, and commercial screening methods. Maintain the Global Deployment Economics and Transferability workbook for selected U.S. and European cases without creating duplicate hardware programs.

Subphase	Principal activity and deliverable	Cost

1A	Dual-reactor utility envelopes, products, site and port assumptions	$12M
1B	31-platform functional architecture and heat-cascade model	$18M
1C	Platform envelope, bay, tunnel, storage and material-flow geometry	$14M
1D	Digital twin, interfaces, hazards, lifecycle cost, and U.S./European transferability model	$14M
1E	Independent reference-design review and frozen procurement baseline	$12M

Work-package hold point

Freeze the architecture only after dual-reactor utility assumptions, site geometry, platform portfolio, hazard boundaries, economic metrics, and interface ownership are independently reviewed.

Area 2. 31 industrial-processing platform designs and pilot prototypes - $180M

Design all 31 platform families and build functional pilots, concentrating hardware and endurance testing on the first five generic thermal platforms, the DAC train, and high-impact sector platforms. The four hydrogen/ammonia rows define standardized platform modules and interfaces, while most ammonia prototype hardware remains funded under Work Package 3 to avoid double counting.

Subphase	Principal activity and deliverable	Cost
2A	FEED for all 31 platforms and first-five concept competitions	$40M
2B	First-five generic platform pilot design/fabrication	$55M
2C	Sector-specific pilots for paper, food, textiles, alumina, chemicals, drying and mineral heat	$45M
2D	Higher-temperature hybrid/electric pilots and controls	$20M
2E	Endurance, teardown, reliability growth and manufacturing packages	$20M

Work-package hold point

Do not scale a platform simply because it has a finished CAD model; require pilot evidence on efficiency, heat integration, maintenance, controllability, and cost before large fabrication awards.

Area 3. Hydrogen and ammonia production platforms - $95M

HTSE/SOEC hydrogen, air separation/nitrogen, Haber-Bosch ammonia, refrigeration/storage/loading, heat recovery, and a repeatable factory-buildable ammonia platform package.

Subphase	Principal activity and deliverable	Cost
3A	HTSE/SOEC + PWR/HTR steam integration and water/oxygen systems	$20M

3B	Nitrogen/air separation, compression and buffer storage	$12M
3C	Haber-Bosch synthesis, recycle, heat recovery and controls	$20M
3D	Integrated ammonia pilot, refrigeration, storage and loading	$30M
3E	Long-duration runs, catalyst/material inspection and scale-up package	$13M

Work-package hold point

Authorize the integrated ammonia pilot only after hydrogen/nitrogen quality, hazardous-area design, synthesis-loop control, storage/loading, relief/containment, and emergency systems pass review.

Area 4. Marine ammonia fuel-cell propulsion prototypes and endurance testing - $70M

Competing direct-SOFC, partial-cracking SOFC, and cracked-ammonia PCFC systems; small integrated propulsion prototypes; thousands-of-hours durability testing and teardown inspection.

Subphase	Principal activity and deliverable	Cost
4A	Direct SOFC, partial-cracking SOFC and PCFC reference designs	$10M
4B	Stack/cracker/catalyst/material bench durability rigs	$15M
4C	100–500 kW integrated propulsion prototypes and emissions controls	$20M
4D	One MW-class shared-load demonstrator, motors/inverters and heat recovery	$12M
4E	5,000–10,000-hour cumulative campaigns, teardowns and marine package	$13M

Work-package hold point

Advance a marine architecture only after stack/cracker degradation, NH3 slip, NOx/N2O control, thermal cycling, marine motion, maintainability, and replacement cost are supported by endurance data.

Area 5. Material handling, tunnels, storage, waterfront, and modular logistics equipment - $85M

Pre-equipped tunnels, pipe/manifold modules, pumps/compressors, conveyors, transfer towers, storage modules, waterfront loading/unloading, and the common distribution spine.

Subphase	Principal activity and deliverable	Cost
5A	Distribution spine, tunnels, manifolds and material-flow design	$15M
5B	Pipe/pump/compressor and utility modules	$15M

5C	Conveyor/transfer-tower and rail-cart modules	$15M
5D	Waterfront loading/unloading and modular storage prototypes	$25M
5E	Integrated material-flow demonstrations and maintenance/recovery tests	$15M

Work-package hold point

Proceed to integrated logistics only after flows, contamination controls, storage inventories, fire/explosion zones, recovery methods, and waterfront constraints are verified.

Area 6. Heavy module transport, double rail, ship/shore transfer, and gantry systems - $65M

Rigid truss, synchronized carriers, 12×24 through 12×96 m module transport, temporary track, loading, ship transfer, alignment, recovery, and repeated handling demonstrations.

Subphase	Principal activity and deliverable	Cost
6A	Load classes, route/port/ship interfaces, multi-rail geometry	$10M
6B	Rigid truss, synchronized carrier and active-jacking design	$15M
6C	1:10 transport/gantry/switch prototypes and controls	$15M
6D	Representative loaded transport and ship/shore transfer	$18M
6E	Reliability, recovery, inspection and final transport standard package	$7M

Work-package hold point

Proceed to loaded full-scale transfers only after loss-of-synchronization, braking, structural deflection, jacking, curve/grade behavior, ship stability, and emergency recovery are demonstrated.

Area 7. Electrified rail and high-power construction-machine interface - $40M

Segmented LV/MV enclosed conductor rail, collectors, switchgear, grounding, interlocks, regenerative operation, warning/lockout systems, and a full-voltage reduced-length test loop.

Subphase	Principal activity and deliverable	Cost
7A	Voltage classes, conductor geometry, collector and grounding requirements	$7M
7B	Segment switchgear, interlocks, detection and lockout architecture	$8M
7C	Full-voltage reduced-length test loop and mobile collector prototypes	$15M

7D	High-power machine, fault, water/dirt and regenerative tests	$7M
7E	Open electrified-rail standard and deployment package	$3M

Work-package hold point

Energize high-power tests only after touch/step potential, insulation failure, water ingress, collector loss, segment isolation, lockout/tagout, fault clearing, and emergency shutdown are independently accepted.

Area 8. Automated construction of thermal-storage tanks - $50M

Tank-jacking, shell-panel placement, automated welding/NDE, floor/roof/insulation handling, modular pump/manifold equipment, and representative tank construction demonstrations.

Subphase	Principal activity and deliverable	Cost
8A	Tank sizes, media, foundation, codes and construction sequence	$7M
8B	Jacking, shell-panel, floor/roof and insulation machine designs	$12M
8C	Automated welding, NDE and dimensional-control cells	$12M
8D	Representative 1:5–1:10 tank build plus full-size weld trials	$14M
8E	Reliability growth, construction-rate model and final machine package	$5M

Work-package hold point

Proceed to representative tank construction only after jacking stability, weld procedure, NDE acceptance, dimensional tolerances, foundation behavior, and safe machine recovery are demonstrated.

Area 9. Standards, interoperability laboratory, digital thread, and conformance testing - $55M

Open mechanical, dimensional, heat, piping, electrical, rail, controls/data, cybersecurity, waterfront, storage, and digital-thread standards plus multi-vendor interoperability tests.

Subphase	Principal activity and deliverable	Cost
9A	Standards steering group and master architecture	$8M
9B	Mechanical/module/rail/piping/thermal/electrical interface drafts	$12M
9C	Controls/data/cyber/digital-thread schemas and conformance tools	$10M
9D	Interoperability laboratory hardware and multi-vendor plug tests	$18M

9E	v2.0 consensus packages and ISO/IEC/Chinese/U.S./European coordination	$7M

Work-package hold point

Freeze v1.0 interfaces only after multiple independent suppliers use the same drawings/schemas in prototype work; freeze v2.0 only after interoperability tests expose and close ambiguity.

Area 10. Integrated prototype campus, flow/corrosion loops, and long-duration campaigns - $65M

Non-nuclear test campus with PWR-like and HTR-PM-like heat loops, full-temperature materials/corrosion systems, site choreography, integrated platform campaigns, maintenance, and reliability growth.

Subphase	Principal activity and deliverable	Cost
10A	Pilot campus, utilities, externally heated PWR/HTR loops and safety systems	$18M
10B	Full-size representative piping/valve/material/corrosion loops	$12M
10C	Instrumentation, digital acquisition and test automation	$8M
10D	Integrated platform/logistics/rail/tank campaigns	$12M
10E	12–18 months of endurance operation, maintenance and analysis	$15M

Work-package hold point

Begin long-duration integrated operation only after test hazards are closed, instrumentation is calibrated, coupon/spool inspection plans are frozen, and stop-work authority is clear.

Area 11. Program management, safety, techno-economics, and independent review - $25M

Owner-side systems engineering, financial control, HAZOP/functional safety, cost and emissions models, independent verification, public data release, and Phase-II decisions.

Subphase	Principal activity and deliverable	Cost
11A	Program office, procurement and configuration control	$6M
11B	HAZOP, functional safety, environmental and marine safety oversight	$5M
11C	Techno-economics, lifecycle emissions, logistics and fossil benchmarks	$5M
11D	Independent technical/cost audits and replication	$5M
11E	Final Phase-II decision packages and public data release	$4M

Work-package hold point

Release public performance/cost claims only after independent replication of key data and explicit separation of measured results from scenario assumptions.

Central contingency and risk reserve - $200M

Reserve stage	Maximum	Examples
Requirements / architecture	$25M	Design-data limits, alternative platform geometry, site/port constraints, standards gaps.
Prototype design	$40M	Failed stack, catalyst, heat exchanger, transport, rail, tank, welding, materials or safety tests.
Fabrication / campus	$55M	Supplier failure, long-lead equipment, rework, test-site changes, replacement prototypes.
Integration / endurance	$60M	Damaged hardware, accelerated degradation, additional campaigns, redesign after long-duration failures.
Final readiness	$20M	Additional reliability runs, conformance tests, standards work, model replication and deficiency closure.

Nuclear Energy Cost Assumptions

Prompt to AI: Use the enclosed "cmp_fission_automate_site_construction,v2j.docx" as the cost basis for Automated China PWR and Automated China HTR-PM electricity and heat. Summarize the ranges and central planning values used by this industrial-processing initiative.

Response From AI:

The industrial-processing proposal adopts the enclosed nuclear site-construction initiative values as planning scenarios. They are not demonstrated market prices. They allow the platform program to compare which industrial heat routes are worth testing and to determine how much platform/logistics cost can be tolerated before a fossil incumbent becomes cheaper.

China nuclear case	Build time	Overnight CAPEX ($/$W_e$)	Electricity LCOE ($/$MWh_e$)	Heat LCOE ($/$MWh_{th}$)	Central planning value	Heat-quality implication
Automated China HTR-PM	3–4 yr	$1.70 – $2.10	$32 – $41	~$13 – $17	~$37/$MWh_e$; ~$15/MWh_{th}	~520–567°C steam; strongest for 300–500°C process heat
Automated China PWR	3.5–4.5 yr	$1.45 – $1.75	$27 – $34	~$9 – $11	~$31/$MWh_e$; ~$10/MWh_{th}	<300°C steam; strongest for low/med. heat & heat-pump/MVR electricity
Conventional China PWR	5–7 yr	$2.00 – $2.50	$34 – $45	~$11 – $14	~$40/$MWh_e$; ~$13/MWh_{th}	Reference incumbent nuclear construction

Costs are summarized as follows:

China planning benchmark	Source	Electricity cost	Useful heat cost
New Automated China HTR-PM (LCOE)	Nuclear ("green")	$37/$MWh_e$	$15/$MWh_{th}$
New Automated China PWR (LCOE)	Nuclear ("green")	$31/$MWh_e$	$10/$MWh_{th}$
New coal — ultra-supercritical (LCOE)	Fossil Fuel	$51/$MWh_e$	$21/$MWh_{th}$
Existing coal — marginal cost	Fossil Fuel	$40/$MWh_e$	$18/$MWh_{th}$

Clarification on existing-coal marginal heat. The $18/$MWh_{th}$ planning value is intended to include coal fuel plus variable pollution-control/filtration cost; it is not a fuel-only number. It excludes sunk

plant/boiler capital and fixed O&M. When this document compares new nuclear with an existing coal site, that is the marginal-cost definition being used.

These values explain why the discussion prioritizes PWR electricity plus heat pumps/MVR below roughly 160–200°C, direct PWR heat at intermediate temperatures, and HTR-PM direct heat where the process genuinely needs roughly 300–500°C. The program should test the entire delivered product cost, not only the reactor energy term.

Industrial Processing to Reduce CO_2 Emissions

Prompt to AI: What things can we do with Automated China PWR and HTR-PM, and platforms of chemical processing equipment, to reduce CO_2 emissions where the green nuclear method cost less than the traditional fossil fuel-based method? In other words, where can Automated China PWR/HTR-PM supply **direct nuclear heat, nuclear-powered heat pumps, or nuclear-electric furnace heat** more cheaply than burning coal, natural gas, oil, or other fossil fuels at the industrial process itself?

Response From AI:

I think the biggest opportunities are **low- and medium-temperature industrial heat**, followed by a smaller but potentially important set of high-temperature hybrid processes.

Your the current companion site-construction proposal assumptions are unusually favorable: Automated China PWR at about $31/MWh-electric and $10/MWh-thermal, and Automated HTR-PM at $37/MWh-electric and $15/MWh-thermal. PWR steam is below roughly 300°C, while HTR-PM can supply roughly 520°C-class steam.

Your earlier China fossil-heat calculation put coal at roughly **$16–23/MWh of useful heat** and natural gas at roughly **$59–90/MWh of useful heat**, before some coal-handling, ash and pollution-control costs.

That gives us the basic economics.

Way of supplying process heat	Nuclear energy cost before process equipment	Approx. useful temperature	Versus Chinese coal heat	Versus Chinese gas heat
Direct Automated PWR heat	~$10/$MWh_{th}$	< ~300°C	Likely cheaper	Much cheaper
Direct Automated HTR-PM heat	~$15/$MWh_{th}$	up to ~500°C useful interface range	Competitive / possibly cheaper	Much cheaper
PWR electricity + heat pump, COP 3.5	~$9/MWh useful heat electricity component	typically <~200°C	Potentially much cheaper	Much cheaper
HTR-PM electricity + heat pump, COP 3.5	~$11/MWh	typically <~200°C	Potentially cheaper	Much cheaper
PWR electricity + resistance furnace	~$32/MWh	potentially >1,000°C	Usually more expensive	Potentially cheaper
HTR-PM heat + electric topping	roughly $20–30/MWh energy component, process-dependent	~500–1,000+°C	Uncertain	Potentially much cheaper

The heat-pump result deserves emphasis. IEA uses a representative industrial heat-pump COP of **3.5** and reports that industrial heat pumps in China are already nearly cost-competitive with coal boilers and substantially better than gas boilers using conventional Chinese energy prices. Your assumed $31/MWh nuclear electricity is much cheaper than the industrial electricity prices in that analysis, so the economics should improve substantially.

Program-wide economic screening rule

The program should screen every platform against the same delivered-energy benchmark before committing major prototype money. Using the central China planning values in this document—$21/MWh$_{th}$ for new coal, $18/MWh$_{th}$ for existing-coal marginal heat, $10/MWh$_{th}$ for Automated PWR heat, $15/MWh$_{th}$ for Automated HTR-PM heat, and $31/MWh$_{e}$ for Automated PWR electricity—the energy-only headroom is:

Nuclear route	Energy component before platform CAPEX	Headroom vs new coal ($21/MWh$_{th}$)	Headroom vs existing coal ($18/MWh$_{th}$)	Screening interpretation
Direct Automated PWR heat	$10/MWh$_{th}$	+$11/MWh$_{th}$	+$8/MWh$_{th}$	Strong starting point for low/medium-temperature heat.
Direct Automated HTR-PM heat	$15/MWh$_{th}$	+$6/MWh$_{th}$	+$3/MWh$_{th}$	Use where 300–500°C heat has real process value.
PWR electricity + heat pump, COP 3.5	$8.86/MWh$_{th}$ electricity component	+$12.14/MWh$_{th}$	+$9.14/MWh$_{th}$	Large headroom before heat-pump capital and O&M.
PWR electricity + heat pump, COP 4	$7.75/MWh$_{th}$ electricity component	+$13.25/MWh$_{th}$	+$10.25/MWh$_{th}$	Especially attractive when waste heat is recoverable.
PWR resistance/electric furnace	$31/MWh$_{th}$ electricity component	−$10/MWh$_{th}$	−$13/MWh$_{th}$	Needs process-efficiency gains or a gas/oil incumbent; not an automatic coal-cost win.
50% HTR-PM heat + 50% electric topping	$26/MWh$_{th}$ simplified energy component	−$5/MWh$_{th}$	−$8/MWh$_{th}$	High-temperature option; economics depend strongly on avoided gas/oil and process efficiency.

Positive headroom is not profit; it is the amount available to pay for heat exchangers, piping, heat pumps, platform capital, maintenance, site logistics, financing, and risk before the nuclear route loses

its energy-cost advantage. Each platform business case should publish this headroom together with expected CO_2 reduction, throughput, capital cost, and sensitivity to fossil-fuel prices.

Where I would concentrate

1. Evaporation, boiling, low-pressure steam and hot water — probably the largest easy opportunity

This is my **#1 target**.

Processes currently burn coal or gas merely to make steam or hot water. Instead:

> Automated PWR → steam/intermediate loop → factory-made process platform

or, where there is a suitable low-temperature waste-heat source:

> Automated PWR electricity → industrial heat pump/MVR → steam

Applications include food processing, beverage production, dairies, breweries, sugar/starch plants, pharmaceuticals, fermentation, chemical plants, textiles, pulp and paper.

IEA reports Chinese industry consumes about **38 EJ/year of heat**, with nearly **20% below 200°C**. Light industry, pulp and paper, and chemicals account for more than three-quarters of this sub-200°C heat demand.

That means this isn't a niche market—it is several thousand TWh/year in China alone.

For direct PWR heat, even allowing several dollars/MWh for the isolation heat exchanger, piping and platform equipment, there is a reasonable chance of staying below Chinese coal heat and a very large margin versus natural gas.

2. Industrial heat pumps + mechanical vapor recompression

This may actually be **the strongest economic opportunity of all**, because you aren't simply replacing one heat source with another—you are reducing the amount of purchased energy dramatically.

Suppose the process needs 1 MWh of 120–160°C useful heat.

At COP 3.5:

> 31 / 3.5 = \$8.86/MWh_{th} for PWR electricity.

At COP 4: 31 / 4 = \$7.75/MWh_{th} for PWR electricity.

> The heat-pump equipment, heat exchangers and maintenance still need to be paid for, so $7.75–8.9 is not the final LCOH. But compare it with **$16–23/MWh coal fuel cost** and **$59–90/MWh natural-gas useful heat**. There is enough economic headroom to pay for considerable equipment.

IEA finds commercially available Chinese industrial heat pumps can address substantial demand below 200°C and identifies standardization as a cost problem because systems currently have to be tailored to particular factories.

That fits your platform concept unusually well.

I would create standardized factory-built units such as:

> 25 MW_{th} heat-pump platform
> 50 MW_{th} heat-pump platform
> 100 MW_{th} steam-generation platform
> MVR evaporation platform

with standardized steam, water, electricity and control interfaces.

3. Pulp and paper

This looks particularly strong.

Much of a paper mill's fossil heat goes into:

- drying paper;
- evaporating black liquor;
- pulp cooking;
- hot water;
- low- and medium-pressure steam.

Your earlier document already identifies pulp/paper as an excellent fit for nuclear steam, heat pumps and vapor recompression.

A good architecture would be:

> PWR steam → high-temperature steam loads

then:

> condensate/waste heat → heat pump → lower-pressure steam

and:

evaporators → mechanical vapor recompression

The one qualification is that mills burning **black liquor, bark or other internal biomass residues** may already have very cheap heat, so those aren't the first sites I would pursue.

Target mills that currently purchase coal, natural gas or fuel oil.

4. Food, beverage, sugar and starch processing

This is another excellent application because temperatures are generally modest and the processes are repetitive enough to standardize.

Typical thermal duties include:

- pasteurization;
- sterilization;
- cooking;
- evaporation;
- drying;
- washing;
- hot water;
- steam.

Globally, IEA identifies food processing and other light industries as a major low-temperature opportunity; roughly 75% of heat in food, textiles and machinery-type light industries is below 200°C.

Here I would strongly favor **PWR + heat pumps + MVR**, rather than HTR-PM. Paying for 500°C-capable reactor heat to make 120°C steam isn't necessary.

5. Textile dyeing, washing and drying

This is particularly relevant in China because the processes require huge quantities of hot water and steam but generally don't require extreme temperatures.

A standardized textile-energy platform could contain:

heat recovery from wastewater
→ heat pump
→ hot-water headers
→ steam generation
→ drying air
→ water recycling

Here a heat pump does double duty: recovering energy that otherwise goes down the drain while replacing a gas/coal boiler.

DOE notes that electrified processes such as microwave, RF and other advanced drying techniques can sometimes reduce the amount of thermal energy required rather than merely substituting electricity for combustion.

That creates an additional economic advantage.

6. Chemical-plant distillation and reboilers

This could be enormous.

Chemical plants burn large quantities of fossil fuel to make steam for:

- distillation-column reboilers;
- evaporation;
- stripping;
- solvent recovery;
- reaction heating;
- drying.

Different temperatures are devoted to different uses:

For duties below ~200°C:	PWR electricity + heat pump/MVR is especially attractive.
For roughly 200–300°C:	Direct PWR steam becomes attractive.
For roughly 300–500°C:	HTR-PM direct heat becomes the stronger option.

Your earlier industrial assessment identifies chemicals and petrochemicals as a major selective opportunity for direct steam, nuclear heat pumps and reboilers.

There is an important distinction: we are replacing **fossil fuel burned for heat**, not fossil hydrocarbons that are chemical feedstocks.

So, a polyethylene plant might still consume petroleum-derived feedstock while its fired heaters, boilers and reboilers become nuclear powered.

7. Alumina refining — one of the best heavy-industry targets

I still think alumina belongs very high on the list, but **not aluminum smelting**, under your new definition.

The Bayer process uses substantial steam for:

> bauxite digestion → evaporation → caustic recovery → precipitation-related heating

and digestion commonly occurs in approximately the 150–250°C range. Your Discussion document identifies it as highly compatible with nuclear steam.

That means **Automated PWR** is probably the economic winner.

You don't even need the HTR-PM for much of the refinery.

A factory-built alumina platform could combine:

**PWR steam

1. digestion
2. flash-steam recovery
3. evaporation
4. MVR
5. condensate recovery**

and drive the effective energy cost below simply burning coal.

The later calcination step needs much higher temperature and should be treated separately.

8. Drying — an enormous cross-industry market

"Drying" sounds mundane, but fossil-fired dryers consume enormous amounts of energy.

Examples include:

- food;
- grain;
- paper;
- chemicals;
- minerals;
- ores;
- coal itself;
- ceramics;
- wood;
- textiles;
- coatings.

DOE specifically identifies grain, coating, paper, aggregate and metal-scrap drying among industrial process-heating opportunities and notes that microwave and other volumetric heating techniques can reduce thermal demand by heating material differently from conventional combustion.

Depending on temperature:

<200°C:	PWR heat pump/MVR.
200–300°C:	PWR steam/direct heat.
300–500°C:	HTR-PM.
Higher:	HTR-PM preheat + nuclear-electric topping.

This is exactly the kind of application where multiple standard drying-platform designs could replace hundreds of one-off fossil dryers.

Where HTR-PM becomes particularly valuable: 300–500°C

PWR dominates much of the lower-temperature market because its assumed heat is cheaper.

The economic purpose of the HTR-PM is therefore **not to compete with the PWR at 120°C**.

Use the HTR-PM where you genuinely need its temperature.

Your the current companion site-construction proposal document makes this distinction explicitly: PWR steam is below roughly 300°C while HTR-PM offers roughly 520°C-class steam.

Candidate applications include:

- superheated-steam drying;
- higher-temperature chemical reactors;
- high-temperature distillation/reboiling;
- mineral treatment;
- feedstock/material preheating;
- selected calcination duties;
- high-temperature thermal-fluid loops.

At $15/MWh reactor heat, HTR-PM heat has only a modest energy-cost advantage versus cheap Chinese coal, so **distance and platform cost matter greatly**.

The ideal HTR-PM customer is therefore **co-located**—perhaps hundreds of meters or a few kilometers away, rather than 30 km away.

Against natural gas, however, the economic margin is much larger.

9. Hybrid HTR-PM + electric furnace: potentially important above 500°C

There is another architecture that I think deserves much more attention in your proposal.

Suppose a process requires **900°C**.

Do not necessarily generate electricity from all the reactor heat and then use electricity from room temperature to 900°C.

Instead:

> HTR-PM direct heat → preheat material/gas to ~500°C

then:

> nuclear electricity → resistance/induction/plasma heating → 900°C

This is potentially much cheaper.

For a simplified case where half the useful heat can be supplied directly at $15/MWh and half needs $37/MWh electricity:

> $0.5 \times \$15 + 0.5 \times \$37 = \$26/\mathrm{MWh_{th}}$ before furnace/platform costs.

That is:

1. above the $16–23/MWh fuel-only coal benchmark;
2. but dramatically below the $59–90/MWh Chinese natural-gas heat benchmark.

And electric heating can sometimes be more efficient at actually putting heat into the material than a combustion furnace because it avoids heating large flows of nitrogen-containing flue gas. DOE identifies resistance, induction, microwave, RF, infrared, electric arc and plasma as candidate replacements for combustion process heat and notes potential reductions in energy intensity, heating time and temperature nonuniformity.

That opens another group of applications.

10. Metal reheating and heat treatment

Not steel reduction—the **heat** operations.

Examples:

- reheating billets;
- forging furnaces;
- annealing;
- heat treatment;
- preheating;
- foundry holding furnaces.

These are commonly fossil-fired.

Possible platform:

HTR-PM preheater → induction/resistance furnace → heat recovery → PWR/HTR steam network

DOE specifically identifies hybrid induction plus conventional radiant/convection heating for steel-billet reheating and other electric process-heating technologies.

With $31/MWh PWR electricity, a fully electric furnace is probably **not reliably cheaper than a coal-fired furnace** purely on energy cost.

It can, however, plausibly be cheaper than **natural-gas/LPG/oil heating**, particularly if induction or direct electrical heating cuts energy consumption per tonne.

So, I would pursue gas-fired metal furnaces first.

11. Glass melting

Same argument.

A glass furnace operates far above HTR-PM temperature, so direct reactor heat alone cannot do the job.

But DOE explicitly identifies **direct-contact electric glass melting** among industrial electrification technologies.

Potential architecture:

**HTR-PM → batch/air/material preheat

- PWR electricity → electric glass melter
- furnace exhaust/waste heat → lower-temperature platforms**

This could plausibly beat a **gas-fired** glass furnace if the direct electric melter has substantially higher end-use efficiency.

I would not yet claim it beats cheap Chinese coal without a plant-specific cost model.

12. Ceramics, bricks and some mineral calcination

Similar story:

> HTR-PM supplies the first several hundred degrees

and:

> electric resistance/microwave/plasma supplies the high-temperature increment.

This is attractive where the incumbent fuel is natural gas, LPG or oil.

Against cheap coal, the economics are less certain.

What I would exclude from the "cheaper than fossil heat" first wave

I would not initially emphasize:

- **Cement clinker.** The kiln is extremely hot, and most cement CO_2 comes from limestone chemistry rather than fuel. Nuclear heat can eliminate fuel emissions but not the calcination emissions. Your existing analysis therefore correctly rates complete cement decarbonization poorly as an early low-cost application.
- **Blast-furnace ironmaking.** Coke is doing chemical reduction as well as making heat. Simply replacing furnace heat doesn't solve that problem.
- **Oil refineries burning their own byproduct fuel gas.** The fuel can have a very low opportunity cost, making nuclear economically difficult despite technically suitable process temperatures.
- **Pure resistance heating replacing a cheap coal furnace.** At ~$31–32/MWh of nuclear-electric heat versus ~$16–23/MWh useful coal heat, there isn't an obvious cost win unless electric heating substantially reduces process energy consumption.

My ranking

If the objective is specifically:

> Reduce CO_2 while also lowering production cost by replacing fossil-fired industrial heat

I would rank the opportunities this way:

Rank	Platform / industry	Preferred nuclear system	Chance of beating fossil heat
1	Heat pump + MVR + steam platform	Automated PWR electricity + waste heat	Very high
2	Food/beverage/sugar/starch thermal processing	PWR direct heat + heat pumps	Very high
3	Chemical distillation / reboilers / evaporation	PWR below 300°C; HTR above	Very high vs gas; high/medium vs coal
4	Pulp and paper	PWR heat + heat pumps/MVR	High
5	Textile dyeing/washing/drying	PWR heat + heat pumps	High
6	Alumina refining	PWR steam + MVR	High
7	General industrial drying	PWR/HTR depending temp.	High
8	300–500°C chemical/mineral heat	HTR direct heat	High vs gas; conditional vs coal
9	Metal reheating/heat treatment	HTR preheat + PWR elect. topping	Potentially high vs gas/oil
10	Glass/ceramics/mineral furnaces	HTR preheat + electric furnace	Promising vs gas; uncertain vs coal

The broader architecture I would pursue

This leads me to a somewhat different vision of the nuclear industrial site.

Don't think primarily:

> reactor → electricity → factory.

Think:

```
HTR-PM 500°C heat
   ↓
300–500°C process platforms
   ↓ recovered heat
PWR/HTR steam at 150–300°C
   ↓
distillation / digestion / drying / evaporation
   ↓ recovered 80–150°C heat
industrial heat pump / MVR
   ↓
hot water and lowest-temperature loads
```

with nuclear electricity used only where necessary to **pump heat uphill or provide the final high-temperature increment**.

That cascading arrangement extracts useful work from the same thermal energy several times before rejecting it. It is particularly compatible with your factory-platform concept, whose proposed open interface explicitly covers nuclear heat, electricity, water and interchangeable industrial modules.

I think **this process-heat market may ultimately be economically more compelling than ammonia**. Ammonia requires converting nuclear energy into hydrogen, then nitrogen/ammonia synthesis, compression and refrigeration. By contrast, when an industrial plant simply needs 180°C steam, a PWR can potentially replace a coal/gas boiler almost directly—and when the process is suitable for a heat pump, **$31/MWh nuclear electricity can effectively become roughly $8–10/MWh of useful process heat before heat-pump capital costs**. That is a very difficult number for fossil combustion to compete with.

31 Nuclear-Integrated Processing Platforms

Prompt to AI: If ≥ 30 different chemical processing platforms were available, and you wanted to maximize CO_2 emissions reduction ("decarbonization"), what platforms would you recommend be developed? Assume each platform is 12 meters wide by 96 meters long. Include support for Direct Air Capture (DAC) and ammonia production.

Response From AI:

The result is a **31-platform portfolio**. Rows 1–25 focus on replacing fossil fuel burned for industrial heat. Rows 26–27 support a modular DAC train, and Rows 28–31 support the hydrogen/nitrogen/ammonia chain.

I would also design the 31 as **repeatable unit-operation platforms**, not 31 complete factories. That is more powerful: the same 12 m × 96 m distillation or heat-pump platform could be replicated in many different chemical plants.

Your current economic assumptions strongly favor this strategy. the current companion site-construction proposal uses about $10/MWh-heat for Automated China PWR, $15/MWh-heat for Automated HTR-PM, $31/MWh-electric for PWR, and $37/MWh-electric for HTR-PM. It also distinguishes PWR heat below roughly 300°C from HTR-PM's much more valuable ~500°C-class process heat. Your industrial-platform proposal is explicitly intended to make such equipment repeatedly in factories, transport it to the site, and connect it through standardized isolated nuclear/industrial interfaces.

The opportunity is enormous. IEA estimates that industries dominated by low-temperature heat and steam represent roughly **70% of global industrial energy consumption** and produced nearly **3 Gt/year of direct energy-related CO_2 in 2023**.

My recommended 31-platform portfolio

I would rank them approximately as follows. "PWR+HP" means PWR electricity powering an industrial heat pump (HP) or mechanical-vapor-recompression system.

Plat-form	Standardized platform family	Typical heat range	Preferred nuclear route	Major fossil-fired processes displaced	Decarbonization priority
1	Industrial heat-pump steam generator	80–160°C	PWR+HP	Coal/gas steam boilers across many industries	Extremely high

Plat-form	Standardized platform family	Typical heat range	Preferred nuclear route	Major fossil-fired processes displaced	Decarbonization priority
2	Mechanical vapor recompression (MVR) evaporator	70–160°C	PWR+MVR	Boiler steam used for evaporation/concentration	Extremely high
3	Low/medium-pressure process-steam platform	120–280°C	Direct PWR heat	Coal/gas/oil boilers	Extremely high
4	Distillation/reboiler platform	100–300°C	PWR heat + HP/MVR	Fired boiler steam in chemical/refinery distillation	Extremely high
5	High-temperature process-heat platform	300–500°C	Direct HTR-PM	Gas/coal/oil process heaters	Very high
6	General evaporation/concentration/crystallization platform	80–200°C	PWR+HP/MVR	Fossil steam	Very high
7	General industrial drying platform	100–300°C	PWR heat/HP; HTR at upper end	Gas/coal-fired dryers	Very high
8	Pulp cooking/digester steam platform	150–200°C	PWR heat	Fossil process steam	Very high
9	Paper-machine drying platform	100–180°C	PWR+HP/MVR	Steam-heated paper dryers	Very high
10	Food/beverage thermal-processing platform	60–180°C	PWR+HP	Pasteurization, sterilization, cooking, evaporation	Very high
11	Sugar/starch evaporation and crystallization platform	70–160°C	PWR+MVR	Bagasse/fossil boilers where fossil is used	High
12	Dairy/food spray-drying platform	100–220°C	PWR+HP/direct steam	Gas-fired spray dryers	High
13	Textile wet-processing and drying platform	60–200°C	PWR+HP	Dyeing, washing, finishing, drying boilers	High
14	Bayer alumina-digestion platform	150–250°C	Direct PWR heat	Coal/gas-fired refinery steam	Very high
15	Alumina evaporation/caustic-recovery platform	100–200°C	PWR+MVR	Fossil steam	High
16	Chemical-reactor / thermal-fluid heating platform	200–500°C	PWR or HTR-PM	Fired thermal-fluid heaters and process furnaces	Very high

Plat-form	Standardized platform family	Typical heat range	Preferred nuclear route	Major fossil-fired processes displaced	Decarbonization priority
17	Petroleum crude/vacuum distillation heat platform	250–400°C	HTR-PM + electric topping	Refinery fired heaters	Very high, but integration harder
18	Ethylene/olefin steam-cracker heating platform	500–900°C+	HTR preheat + PWR electric furnace	Natural-gas/fuel-gas cracking furnaces	Enormous potential, technically harder
19	Alumina calcination platform	~850–1,100°C	HTR preheat + electric topping	Gas/coal calciner	High
20	Steel reheating/forging furnace platform	800–1,250°C	HTR preheat + electric/induction	Gas/coal reheating furnaces	Very high
21	Steel annealing/heat-treatment platform	500–1,000°C	HTR + electric topping	Gas-fired heat-treatment furnaces	High
22	Glass-melting platform	~1,200–1,600°C	HTR preheat + electric melting	Gas/oil glass furnaces	High
23	Brick/ceramic kiln platform	~600–1,300°C	HTR preheat + electric kiln	Coal/gas kilns	High
24	Lime/cement calcination-heating platform	~850–1,450°C	HTR preheat + electric heating	Coal/petcoke/gas kiln heat	Enormous, but doesn't eliminate process CO_2
25	Mineral drying/roasting/calcining platform	~200–1,000°C	PWR/HTR + electric topping	Coal/gas mineral furnaces	High and broadly reusable
26	Direct-air-capture air-contactor + solid-sorbent regeneration platform	80–120°C	PWR heat + electricity preferred; HTR-PM possible	Removes atmospheric CO_2 directly; requires durable storage	Extremely high if cost/MRV targets are met
27	DAC CO_2 dehydration, compression and storage-interface platform	Ambient / low-temp	PWR electricity preferred	Conditions captured CO_2 for pipeline, ship or geologic storage	Essential companion to DAC
28	HTSE/SOEC hydrogen-production platform	High-T electrolysis	PWR/HTR steam + nuclear electricity	Replaces fossil H_2 used in conventional ammonia and chemicals	Extremely high for green-ammonia chain
29	Air-separation / nitrogen-production platform	Near ambient	PWR electricity	Supplies N_2 for ammonia without fossil-powered auxiliary electricity	High / enabling

Plat-form	Standardized platform family	Typical heat range	Preferred nuclear route	Major fossil-fired processes displaced	Decarbonization priority
30	Haber-Bosch ammonia synthesis, recycle and heat-recovery platform	350–500°C	Nuclear H_2 + electricity; PWR/HTR heat integration	Replaces conventional fossil-hydrogen ammonia production	Extremely high with green H_2
31	Ammonia refrigeration, storage and loading platform	Low-temp / ambient	PWR electricity	Completes green-ammonia chain for fertilizer, export and marine fuel	High / enabling

DAC and Ammonia

Rows 26–27 make DAC (Direct Air Capture) modular. A large DAC facility would use many replicated air-contact/regeneration trains, followed by common CO_2 dehydration, compression, metering and storage-interface equipment. The PWR is the preferred energy source for low-temperature solid-sorbent DAC because the current planning model gives it both cheaper electricity and cheaper low-temperature heat than HTR-PM.

Rows 28–31 make the ammonia chain explicit in the standardized-platform table rather than leaving ammonia only in a separate work package. They correspond directly to the existing program's HTSE/SOEC hydrogen, nitrogen/air separation, Haber-Bosch synthesis/recycle/heat recovery, and refrigeration/storage/loading functions. Most prototype spending for these four rows remains in Work Package 3, so their appearance in the 31-platform architecture does not double the $1B budget.

I would **not divide the R&D funding equally among these 31**.

The first five matter disproportionately

If the purpose is maximum CO_2 reduction per development dollar, I would put perhaps 50–60% of the process-platform R&D effort into Platforms 1–5.

The reason is that they are generic rather than industry-specific.

#1 Heat-pump steam platform could service food, chemicals, paper, pharmaceuticals, textiles, fermentation and dozens of other industries. Industrial heat pumps can already provide a meaningful portion of low-temperature industrial heat; IEA identifies paper, food and chemicals as the largest near-term opportunities.

#2 MVR platform is potentially even more important for plants that evaporate water. Instead of repeatedly boiling water with fresh steam and throwing away the latent heat, the system compresses

the vapor and uses it again. This can radically reduce the energy purchased per tonne of water evaporated.

#3 PWR steam platform is the simplest concept: take low-cost nuclear heat through an isolated heat exchanger and replace a coal/gas boiler.

#4 Distillation/reboiler platform attacks one of the enormous thermal loads in chemicals and refining.

#5 HTR-PM high-temperature heat platform extends the plug-and-play system into approximately 300–500°C applications where PWR heat can't go.

Your existing Discussion document reaches essentially the same strategic conclusion: the strongest initial sectors are **food, pulp and paper, chemicals and alumina**, and the overall architecture should emphasize heat pumps/MVR below ~160–200°C, nuclear steam at intermediate temperatures, and HTR-PM direct heat around 300–500°C.

Why I would make #1 a heat-pump platform rather than simply a steam platform

This is potentially one of the most economically important aspects of the whole concept.

At your assumed Automated China PWR electricity price of $31/MWh, a heat pump with COP = 3.5 uses:

31 / 3.5 = $8.86/$MWh_{th}$ of electricity.

COP = 4 gives: 31 / 4 = $7.75/$MWh_{th}$ of electricity.

Those aren't complete delivered-heat costs because the heat pump has capital and maintenance costs. But your own nuclear model puts existing Chinese coal heat around **$18/MWh** and new coal around **$21/MWh**.

That leaves substantial room to pay for the heat pump.

More importantly, the machine isn't creating all of the heat from nuclear electricity. It is **recycling heat that the factory previously discarded**. That is why it can displace several units of fossil heat for every unit of nuclear electricity consumed.

IEA notes that about three-quarters of the heat demand in food, textiles and similar lighter industries is below 200°C and therefore potentially addressable with existing low-emission heat technologies.

Platforms 2, 4 and 6 could be an especially powerful family

I'd actually design these together.

A chemical plant frequently has:

> distillation → condensation → evaporation → cooling

at different temperatures.

Today it may burn gas or coal to create reboiler steam at one location while simultaneously throwing heat into a cooling tower somewhere else.

The platform system should instead do:

> hot condenser/waste heat
> → heat pump
> → reboiler

and:

> evaporated process vapor
> → MVR compressor
> → same evaporator heating surface

with nuclear steam supplying only the temperature/energy deficit.

That could make nuclear industrial decarbonization much cheaper than simply replacing every fossil boiler with a nuclear boiler on a one-for-one basis.

Alumina deserves two or three dedicated platforms

I would put alumina unusually high in the program because it combines **huge industrial scale with favorable temperatures**.

The existing Discussion document identifies alumina as a strong early demonstration because digestion, evaporation and caustic recovery can use nuclear steam and heat recovery, while higher-temperature calcination can be separately electrified or hybridized.

Rather than try to fit an alumina refinery onto one platform, I would use:

Platform 14: bauxite slurry + digestion heat exchangers.

Platform 15: evaporation + caustic concentration + MVR.

Platform 19: calcination.

Multiple identical platforms would operate in parallel for a large refinery.

That is exactly the advantage of fixing the transportation envelope at 12 m × 96 m: **capacity is obtained by replication instead of building ever-larger one-off equipment.**

Chemical plants may be the single largest customer

Platforms **4, 6, 16, 17 and 18** form a temperature ladder for chemical and petrochemical processing:

<200°C:	Heat pumps, MVR, evaporation.
150–300°C:	PWR steam and reboilers.
300–500°C:	HTR-PM heat-transfer fluid.
500–900°C:	HTR-PM preheat + electric topping.
> ~900–1,000°C:	primarily electric furnace, with HTR-PM supplying whatever preheat can economically be recovered.

This matters because many chemical processes require fossil fuels both as **feedstocks** and as **heat sources**. We're deliberately addressing only the second category here.

For example, an ethylene steam cracker may still consume hydrocarbon feedstock. But its enormous gas-fired furnace can potentially be replaced by an electrically heated cracking furnace, with HTR-PM supplying feed preheat and recovered heat being cascaded into other platforms. IEA notes that conventional steam cracking requires temperatures near **1,000°C**, illustrating why this belongs in the harder high-temperature tier rather than among the initial PWR steam applications.

Cement is different

I would still develop Platform #24 because cement is so large globally, but I wouldn't let it dominate the first-generation program.

A kiln can potentially eliminate much of its **fuel-combustion CO_2** by replacing coal/petcoke heat with nuclear-derived heat/electric heating.

But limestone undergoes:

$$CaCO_3 \rightarrow CaO + CO_2$$

so nuclear heat cannot eliminate that process CO_2.

High-temperature heat represents a very large industrial energy load, but IEA estimates industrial process emissions themselves account for roughly one-quarter of industrial emissions; fuel substitution doesn't address them.

Thus, cement gives large absolute reductions, but less CO_2 reduction per tonne of product than one might initially expect from replacing the kiln fuel alone.

An important 26th item: not really a chemical platform

In addition to the 31 processing platforms, I would still develop one common utility module outside the numbered portfolio:

Nuclear Industrial Heat Hub

It would contain the standardized boundary between reactor and chemical yard:

PWR steam
HTR-PM heat-transfer loop
electricity
cooling water
demineralized water
condensate return
thermal storage
meters
isolation valves
emergency shutdown
controls

Then every one of the 31 platforms plugs into the same headers.

That architecture is already implicit in your $1B platform proposal: the fund calls for an open nuclear-to-process interface with an intermediate isolation loop so industrial equipment isn't directly connected to the reactor primary system.

The physical site I envision

The interesting result is that these platforms shouldn't be scattered randomly around the nuclear plant.

I would arrange them **by decreasing temperature**, allowing heat to cascade:

HTR-PM ~500°C header
→ chemical reactors / high-temperature preheaters
→ **300°C heat**
→ distillation / alumina / drying
→ **150–200°C heat**
→ evaporation / steam / paper / food
→ **60–120°C heat**
→ heat pumps
→ hot water / lowest-temperature processing
→ cooling system.

Meanwhile, the Automated PWR supplies tremendous quantities of the cheapest low-temperature nuclear heat and cheap electricity to the heat pumps, MVR compressors and high-temperature electric topping heaters.

That is much better thermodynamically than connecting each platform independently to the reactor and throwing away its exhaust heat.

Therefore, if I had to prioritize the expanded portfolio

I would spend the development effort heavily toward generic heat manipulation rather than treating all 31 platform families as equally large, product-specific development programs.

My first five would be:

- **Industrial heat-pump steam generator**
- **MVR evaporator**
- **PWR process-steam platform**
- **Distillation/reboiler platform**
- **HTR-PM 300–500°C heat platform**

Those five could probably be deployed in more factories and displace more fossil-fired process heat than the other 20 combined initially. Then the remaining 26 platforms adapt these basic thermal technologies to large sector-specific loads—paper, food, textiles, alumina, refining, petrochemicals, steel, glass, ceramics and cement.

That is also why I would broaden the name from **“Chemical Processing Platforms”** to something like **“Nuclear-Integrated Industrial Processing Platforms.”** Much of the largest decarbonization opportunity is not making a new molecule; it is replacing the enormous amount of coal and gas presently burned to **boil, evaporate, dry, distill, heat, roast and calcine material**.

Material Handling Costs

Prompt to AI: What kind of costs are associated with getting materials into and out of the chemical processing site (to/from ships)? How do these costs compare to the traditional fossil fuel-based approach? Also, there is two ways to look a this: (a) new construction at nuclear site vs new construction at tradition fossil fuel site, (b) new construction at nuclear site vs margin cost of existing fossil fuel site. Note that existing fossil fuel site margin cost is cost of fuel plus cost of filtration. Assume we are in China and working with Coal.

Response From AI:

Yes. The main conclusion is that material-handling cost is important, but it is usually much smaller than the process-heat advantage created by the assumed Automated PWR/HTR-PM costs—provided the chemical yard is built at a deep-water site and the process chain is kept together.

Your concept already points in the right direction: 12 m × 96 m platforms are clustered next to the reactors, with shared underground material/utility corridors and a common waterfront, while entire platforms can be brought in and replaced through the double-rail system. The platform-development proposal also explicitly includes port transfer, marine transport, common storage, product export, and reusable utility systems.

I would distinguish three different logistics costs:

- **Ship ↔ site**: loading/unloading, berth, pumps/conveyors, temporary storage.
- **Within the site**: pipes, conveyors, rail carts, tunnels, manifolds, pumping/compression.
- **Moving intermediates off-site and back again**: this is the one I would try hardest to eliminate.

1. What does ship-to-site material handling cost?

Chinese port data give a useful reality check. A Tangshan import-coal tariff listed about **RMB 27.5/t**, roughly **$4/t**, for a package including ship unloading, weighing and internal port handling. A 2024 Qingdao Port filing gives liquid-bulk port handling around **RMB 15.09/t (~$2.2/t)** and tank storage around RMB 9.4–13/t; longer pipeline transportation to customers is substantially more expensive, which is exactly what co-location avoids.

For your highly automated captive terminal, I would therefore use these planning numbers:

Material movement	New integrated China site: planning cost
Liquid transferred ship ↔ nearby platform by pipe	$2–5/t

Gas/liquefied gas / hazardous liquid	$3–7/t
High-volume dry bulk by enclosed conveyor	$3–6/t
Less convenient dry bulk	$4–8/t
Packaged/drummed/palletized solids by rail cart	$6–15/t
Platform-to-platform pumped liquid, short distance	~$0.2–1/t
Platform-to-platform bulk conveyor	~$0.3–1.5/t
Platform-to-platform rail-cart movement	~$1–3/t

These last several numbers are my **engineering planning estimates**, rather than published Chinese tariffs.

The actual energy required to pump a liquid several hundred meters is almost negligible. The money is mainly in **pipes, pumps, valves, manifolds, maintenance, inspections, containment and safety**, not electricity.

2. The new nuclear site needs significant logistics infrastructure

Your concept isn't free. A new nuclear-processing super-site needs:

> berths + loading arms + tanks/silos + pumps + enclosed conveyors + material tunnels + utility tunnels + rail-cart system + double-rail platform transporter + manifolds + safety systems + controls.

Chinese projects show how wide the possible capital range is. A relatively modest liquid-bulk terminal at Tongling is about RMB 100 million for 1.1 Mt/y capacity, whereas the much larger and more sophisticated Huanghua liquid-chemical terminal—two berths plus **54 storage tanks**—is about RMB 4.214 billion for 4.27 Mt/y (i.e. “Mt” refers to million of tonnes, “y” refers to years).

At a 7% real capital cost over 30 years, those examples imply roughly **$1/t to $12/t of throughput just for capital recovery**, before operating cost. The huge range comes from how much storage, hazardous-material capability, dredging, berth work and ancillary equipment is included.

For your very large super-site, I would aim for:

> ~$3–7 per tonne crossing the ship/site boundary, all-in, for high-volume ordinary bulk liquids/solids.

More difficult hazardous chemicals could be **$5–12/t**.

The scale of your facility helps enormously because the same terminal, tunnels, manifolds, storage and automation can support many platforms.

3. New nuclear versus a new coal-based chemical site

This is the easiest comparison.

Both sites need essentially the same logistics for the **chemical feedstocks and finished products**. If we are making alumina, both need bauxite handling. If making petrochemicals, both need hydrocarbon feedstock handling. If exporting methanol, both need methanol tanks and loading equipment.

So those costs largely **cancel**.

But the new coal plant has additional logistics that the nuclear plant doesn't:

> coal ship unloading → coal storage → coal reclaim → conveyors → crushing/milling → furnace → ash/slag handling, plus pollution-control reagents and waste.

China's recent Huanghua coal-terminal expansion is a good indication of scale: roughly RMB 4.9 billion adds about 53 Mt/y of coal throughput. China is able to automate this extremely well—Huanghua already uses automated coal unloading, stacking, reclaiming and conveying.

But it still costs money.

A useful rule is that 5,500-kcal/kg coal contains about 6.4 MWh_{th}/t. At ~85% useful industrial-boiler efficiency: 6.4 × 0.85 ≈ 5.44 MWh_{th} of useful heat per tonne of coal.

So, a ~$4/t coal port-handling cost represents only about $4 / 5.44 ≈ $0.74/$MWh_{th}$ of useful heat.

Add coal-yard/conveyor operation and ash handling and I would carry roughly:

> $1–2/$MWh_{th}$ of additional coal logistics

for a large coastal facility.

That isn't enormous—but nuclear avoids it entirely.

4. Now compare the actual economics

Your the current companion site-construction proposal document currently uses:

Heat source, China	Cost of useful process heat
Automated PWR	$10/$MWh_{th}$
Automated HTR-PM	$15/$MWh_{th}$

New coal	$21/MWh$_{th}$
Existing coal, marginal	$18/MWh$_{th}$

The nuclear values are R&D target/scenario estimates, not demonstrated costs. That's an important qualification.

Case A — New nuclear site versus new coal site

The energy advantage is therefore:

PWR: $21 - $10 = $11/MWh$_{th}$ lower than new coal.

HTR-PM: $21 - $15 = $6/MWh$_{th}$ lower than new coal.

Now consider a chemical product requiring **3 MWh of process heat per tonne of product**.

New-build comparison	Heat cost / ton product	Relative difference
Automated PWR	$30/t	$33/t cheaper than coal
Automated HTR-PM	$45/t	$18/t cheaper than coal
New coal	$63/t	baseline

On top of that, coal has roughly another few dollars per tonne of product for coal handling/ash logistics.

The nuclear yard may have somewhat more expensive plug-and-play tunnels and modular platform interfaces than a conventional chemical plant, but **there is plenty of economic headroom** in a 3-MWh/t process.

I would expect:

- New PWR processing site versus new coal site: nuclear strongly favored.
- New HTR-PM processing site versus new coal site: nuclear favored when its higher-temperature heat is genuinely useful.

Case B — New nuclear versus an existing coal site

This is much harder and is the comparison that should receive the most attention.

For the existing coal plant, the furnace, coal yard, conveyors, buildings, port equipment and pollution-control hardware are already built. Under your definition, we count essentially:

Marginal-cost definition. The \$18/MWh_{th} existing-coal benchmark used below is not fuel only. It includes coal fuel plus variable pollution-control/filtration cost, while excluding sunk capital and fixed plant costs.

> coal fuel + variable pollution-control/filtration cost

rather than recovering sunk capital.

Your the current companion site-construction proposal central value is about \$18/MWh_{th} for that situation.

Therefore:

> Automated PWR advantage: \$18 - \$10 = \$8/MWh_{th}.
>
> Automated HTR-PM advantage: \$18 - \$15 = \$3/MWh_{th}.

Using our same 3-MWh/t chemical:

Comparison against existing coal	Process-heat cost	Heat-cost advantage
New Automated PWR	\$30/t	saves \$24/t product
New Automated HTR-PM	\$45/t	saves \$9/t product
Existing coal, strict marginal	\$54/t	baseline

But now the new nuclear complex must recover its **new waterfront, tunnels, manifolds, storage and other common infrastructure**, whereas much of the existing site's infrastructure is sunk.

I would provisionally allow roughly **\$4–8/t of product** as the additional capital-recovery burden of the new integrated nuclear logistics system relative to an existing site, for a large high-throughput chemical complex.

That gives roughly:

Example: 3 MWh_{th}/t product	Approx. advantage after extra new-site logistics
PWR vs existing coal	~\$16–20/t cheaper
HTR-PM vs existing coal	~\$1–5/t cheaper

This is why I am much more confident that an Automated PWR could beat the marginal-cost incumbent than that an HTR-PM necessarily would.

For HTR-PM, the customer should actually **need 300–500°C heat**. Otherwise PWR is the economically stronger heat source.

A useful break-even rule

Suppose construction of the new waterfront/platform infrastructure gives nuclear an additional disadvantage of **$4–8/t of finished product** relative to an already-built coal plant.

The heat intensity needed to recover that cost is:

New nuclear vs existing coal	$4/t logistics disadvantage	$8/t logistics disadvantage
PWR — saves $8/$MWh_{th}$	0.5 MWh_{th}/t product	1.0 MWh_{th}/t
HTR-PM — saves $3/$MWh_{th}$	1.3 MWh_{th}/t	2.7 MWh_{th}/t

This is an extremely useful screening criterion.

A process consuming **2–5 MWh_{th} per tonne** looks quite promising for a new nuclear super-site.

A process consuming only **0.2 MWh/t** probably doesn't justify abandoning a functioning fossil site solely to obtain cheaper heat.

The much bigger opportunity: don't ship intermediate products away

I think your comment about doing as much processing as possible at the nuclear site is **economically very important**.

Suppose Platform A makes an intermediate chemical.

If you send it elsewhere, process it, and return it, you may incur:

> Site A loading → Site B unloading → Site B loading → Site A unloading.

At only **$3–6/t per terminal movement**, that becomes:

4 terminal movements × $3–6/t = $12–24/t before paying any ocean freight at all.

Add storage, inventory financing, contamination/loss, paperwork and transport, and an out-and-back intermediate-product excursion could easily cost:

$15–35/t of intermediate material, plus the actual ship voyage.

That can be considerably larger than the difference between nuclear and fossil heat.

So I think a central design principle for the Industrial Processing Yard should be:

> Raw material comes into the waterfront once. Keep it inside the site as it moves through as many processing stages as economically practical. Export the highest-value final product—not intermediate products that need to come back.

That is exactly where your underground pipes/conveyors/rail carts become economically interesting.

Moving something **300 meters from Platform A to Platform B** can cost well under $1/t in a high-volume automated system.

Moving the same material:

> Platform A at nuclear site X → storage → ship → another terminal (dock) → factory → terminal (dock) → ship → Platform B at nuclear site Y

can cost tens of dollars per tonne.

I would therefore design the yard somewhat differently

I would make the transportation system hierarchical:

Material	Preferred movement
Very-high-volume liquids	Permanent pipes/manifolds
Very-high-volume gases	Permanent pipe headers
Ore, coal-like material, limestone, salt	Enclosed belt conveyors
Powders requiring enclosure	Pneumatic/screw conveyor or sealed rail carts
Drums, catalysts, bags, pallets	Automated rail carts
Intermediate liquid chemicals	Platform-to-platform pipe network
Entire 12 × 96 m processing module	Double-rail platform transporter

The giant double-rail transporter should **not** ordinarily carry process material. Its job is moving entire platforms during installation, major overhaul and replacement. Your platform-development proposal makes the same basic distinction: ordinary product handling is through utilities/storage/export systems, while the large rigid-truss transport system exists to move the platform itself.

Bottom line

For new construction versus new construction, I don't see material logistics as a significant economic obstacle to the nuclear concept. Both plants need docks and product/feedstock handling; coal additionally needs a fuel-handling and waste-handling system. With your assumed heat costs, the

energy advantage—\$11/MWh_{th} for PWR and \$6/MWh_{th} for HTR-PM versus new coal—is considerably larger than the likely logistics differential.

For new nuclear versus an existing coal plant's marginal cost, logistics matter much more because the existing infrastructure is sunk. At \$10 versus \$18/MWh_{th}, the Automated PWR still has a large \$8/MWh heat advantage and could plausibly overcome the cost of constructing the new logistics network. At \$15 versus \$18/MWh_{th}, the HTR-PM has only \$3/MWh of headroom, so I would use it primarily where its ~500°C heat produces a genuine process advantage.

And I think the largest logistics savings may come from site integration rather than nuclear itself: arranging 31 or more different processing-platform types so that raw materials enter the site once, intermediates move cheaply from platform to platform, waste heat and byproducts are reused locally, and only final products leave by ship. That could save tens of dollars per tonne in processes that otherwise move intermediate products between separate industrial sites.

Large Material Handling Equipment

Prompt to AI: If 12m x 96m platforms of chemical processing equipment are made at a factory or shipyard, and transported to a nuclear site, then what else that is large could be transported to the site to help with material handling, transport, storage, and ship loading/unloading?

Response From AI:

Yes. I would extend the same idea beyond the **chemical-processing platform** and create a family of **large factory-made logistics modules** that can be moved by the same special ship + extra-wide double-rail transportation system.

That is consistent with the architecture in your business plan: pages 54–56 envision 12 m × 96 m processing platforms, double-rail transportation, specially built ships, underground tunnels, and dedicated R&D for the platform transportation system and chemical-processing yard.

The key idea would be:

> Do not merely factory-build the chemical plant. Factory-build as much of the port, storage, material-handling, utility, and interconnection infrastructure as practical.

This could substantially reduce site construction and make each new nuclear industrial complex more like assembling standardized blocks.

Large modules I would transport to the site

Factory-made module	Typical transported form	What it does	Value of factory construction
Platform 1. Pipe-rack / manifold	12 × 24 m or 12 × 96 m	Large pipes for liquids, gases, steam, condensate, ammonia, acids, etc.	Probably one of the highest-value standardized modules
Platform 2. Pumping station	12 × 24–48 m	Pumps, motors, valves, filters, meters	Arrives fully piped/wired/tested
Platform 3. Gas compressor	12 × 24–48 m	Compresses H_2, N_2, NH_3 vapor, CO_2, process gases	Expensive precision equipment benefits greatly from factory assembly
Platform 4. Liquid-storage tank module	Complete smaller tanks; modular sections for huge tanks	Intermediate/product storage	Reduces field welding and installation
Platform 5. Pressurized-storage vessel	Complete horizontal / vertical vessels	Ammonia, LPG-like products, compressed gases	Excellent candidate for complete factory fabrication

Platform 6. Bulk-solids silo modules	Large factory-made bins or sectional silos	Salt, limestone, ore, fertilizer, powders	Standardized discharge and weighing interfaces
Platform 7. Enclosed conveyor modules	~24–96 m conveyor sections	Bulk solids between ship, storage and platforms	Very repetitive equipment; ideal for modular manufacture
Platform 8. Conveyor transfer-tower	12 × 12–24 m tall module	Changes direction/elevation; screens, crushers, samplers	Complex part of conveyor system can arrive preassembled
Platform 9. Rail-cart material station	12 × 24 m	Loads/unloads pallets, drums, catalyst, bags, containers	Automates small-volume solids movement
Platform 10. Ship liquid / gas manifold module	Large quay-side skid/platform	Connects ships to shore pipelines	Standard loading arms, ESD valves, metering and vapor-return system
Platform 11. Bulk ship-loader / unloader module	Large truss/gantry sections	Moves ore, grain, salt, fertilizer, limestone, etc. ship ↔ conveyor	Much of structure can be built at shipyard
Platform 12. Ship-to-rail / platform transfer station	Large waterfront module	Transfers entire 12 × 96 m platform ship ↔ transporter	Essential to your platform-replacement concept
Platform 13. Storage-tank pump-house	12 × 24–48 m	Pumps and manifolds connecting tank farm to yard	Standardized, replaceable unit
Platform 14. Product blending / loading	12 × 24–48 m	Blending, sampling, metering prior to export	Keeps finished-product logistics automated
Platform 15. Intermediate - material distribution	12 × 24–48 m	Central manifold that routes product from Platform A to B/C/D	Reduces duplicate pipes and valves
Platform 16. Utility-distribution	12 × 24–96 m	Steam, hot water, cooling water, compressed air, nitrogen	Common interface for many processing platforms
Platform 17. Heat-exchanger station	12 × 24–48 m	Nuclear isolation loop ↔ industrial heat network	Particularly important for PWR/HTR-PM integration
Platform 18. Electrical substation/pwr-conversion	12 × 24–48 m	Transformers, switchgear, rectifiers, variable-speed drives	Almost completely factory-testable
Platform 19. Heat-pump/MVR utility	12 × 24–96 m	Recycles waste heat and produces process steam	Serves many adjacent chemical platforms
Platform 20. Water-treatment	12 × 24–96 m	Demineralized/process water, condensate polishing	Common site utility
Platform 21. Wastewater-treatment	Multiple 12 × 96 m modules	Treats process water before reuse/discharge	Can be replicated as capacity grows
Platform 22. Air-separation	Multiple large modules	O_2/N_2/Ar production	Useful across NH_3, metals and chemical processes
Platform 23. Refrigeration	12 × 24–96 m	NH_3/product refrigeration, cold storage	Factory-built compressors/heat

			exchangers
Platform 24. Fire-protection/safety	12 × 24–48 m	Fire water, foam, emergency pumps, gas detection controls	Standard site-wide safety package
Platform 25. Central control/automation	12 × 24–48 m	Controls logistics, manifolds, conveyors, carts and storage	Enables largely autonomous material movement

I would go one step further with the underground tunnels

Your page 55 concept already proposes factory manufacture of **precast concrete tunnel modules** or factory-built rebar/form assemblies that are transported to the site.

I think this should be pushed hard.

Instead of just transporting empty concrete tunnel sections, some could arrive as **pre-equipped tunnel modules**.

For example, imagine a 12 m × 24 m precast tunnel section arriving with:

- pipe supports already installed;
- perhaps several large stainless/carbon-steel pipes already mounted;
- cable trays;
- ventilation ducts;
- lighting;
- drainage;
- fire detection;
- instrumentation;
- rail-cart track;
- conveyor support structure.

The module gets lowered into an excavated trench and connected to the next one.

That could turn a great deal of traditionally slow construction into:

> factory → ship → double rail → excavation → connect.

Your business plan already proposes using factory-made tunnel structures; this simply extends the idea by installing much of their internal equipment before shipping.

Storage deserves special treatment

I would **not** necessarily transport every large storage tank intact.

A 100,000 m³ atmospheric tank contains mostly empty space during transportation. Shipping the entire empty tank wastes transportation volume.

So, I would divide storage into two categories.

For pressure vessels and moderate-size tanks, transport the **complete vessel**, already tested.

For enormous atmospheric tanks, transport factory-made **tank kits**:

> floor sections + curved wall panels + roof sections + structural rings + piping/manifold module + pump module.

A specialized automated assembly machine at the site could join these.

That is similar to your underlying philosophy for nuclear construction: **maximize factory work, but don't transport enormous volumes of air if field assembly of large prefabricated pieces is cheaper.**

Conveyors are an especially attractive target

For solids, I would make the conveyor system almost completely modular.

Suppose the waterfront is 1 km from a particular processing area.

Instead of constructing a custom conveyor on site, manufacture perhaps:

> 10 × 96 m enclosed conveyor modules

plus several short connection sections.

Each could contain:

1. belt or screw conveyor;
2. structural truss;
3. weather enclosure;
4. dust collection;
5. fire detection;
6. sensors;
7. motors;
8. controls;
9. maintenance walkway.

They arrive on the same transportation infrastructure used for processing platforms.

The crane places them onto standardized supports.

That could make a kilometer-long bulk-material system much closer to assembling a kit than constructing an industrial facility.

The ship terminal itself could become modular

I think this is another important extension of the concept.

The waterfront could have standardized **berth modules**.

For a liquid/gas berth:

```
ship loading arms
   → emergency shutdown valves
   → vapor return
   → pumps
   → meters
   → sampling
   → pipe manifold
```

could largely reside on one factory-built module.

For bulk solids:

```
ship unloader
   → receiving hopper
   → crusher/screen if needed
   → dust control
   → conveyor
```

could be assembled into several giant shipyard-built modules.

This is particularly logical because shipyards already specialize in constructing enormous steel structures immediately adjacent to water.

You could even build the waterfront equipment at the same shipyard that builds the 12 × 96 m chemical platforms.

I would build a common "distribution spine"

The physical arrangement I favor would have a large central material-distribution corridor running the entire length of the chemical yard.

Think of it almost as an industrial subway.

Above/below it would be standardized networks for:

liquids/gases:	large header pipes;
bulk solids:	enclosed conveyors;
discrete solids:	automated rail carts;
utilities:	steam, water, power and communications;
platform replacement:	the much larger extra-wide double rail.

Each 12 × 96 m processing bay would connect perpendicular to this spine.

Then a processing platform doesn't need a dedicated 800-meter pipe back to the port.

It connects perhaps 10–30 meters to the nearest standard header.

That substantially reduces incremental cost when another platform is installed.

Central storage can also eliminate repeated handling

This becomes particularly important if you want to perform multiple processing steps at the same site.

Suppose:

> Ship → Raw Material A → Platform 1 → Intermediate B → Platform 2 →
> Intermediate C → Platform 3 → Finished Product D → Ship

I would not install separate large storage facilities between every platform.

Instead, through the common distribution spine, there might be:

- Raw-material storage zone
- Intermediate-product tank/silo farm
- Finished-product storage zone

That lets Platform 1 send intermediate B either directly to Platform 2 or temporarily into a common storage tank.

Consequently, platforms can stop independently without shutting down the whole chain.

The transporter itself can move more than chemical platforms

This is an important implication of the double-width rail system described in your plan.

The business plan already envisages 12 m × 24 m double-width railcars transporting very large components and construction material, while the chemical-processing section extends that into four-car 12 × 96 m platform transport.

Therefore, I wouldn't define the system as a **chemical-platform transporter**.

I'd define it as a:

> Heavy Industrial Module Transportation Standard

For example:

> 1 railcar = 12 × 24 m module
> 2 railcars = 12 × 48 m module
> 3 railcars = 12 × 72 m module
> 4 railcars = 12 × 96 m module

Then factories can build equipment to standardized envelopes.

A 96-m ammonia platform uses four.

A compressor station might use one.

A pump station might use one.

A giant air-separation unit might use two.

A conveyor section might use four.

A tunnel section might use one.

A large pressure vessel might use one or two.

This is analogous to the standard shipping container, but **hundreds of times larger**.

The biggest opportunity may be standardized interfaces

The physical module isn't necessarily the hard part.

The major economic benefit comes if **every module plugs into the same standard interfaces**.

For instance, every 12 × 96 m bay could have standardized locations for:

A — high-temperature heat supply
B — heat return
C — medium-pressure steam
D — condensate
E — cooling water supply/return
F — process water
G — high-voltage electricity
H — communications/control
I/J/K — pumped material headers
L — bulk-solids conveyor interface
M — rail-cart interface

The platform transporter arrives, lowers the platform into position, and automated or semi-automated connectors make these interfaces.

Then you can eventually do something very powerful:

Replace not merely individual machines, but entire factories.

If Platform #37 is a 20-year-old evaporator system and a newer factory makes one that's 15% more efficient, disconnect #37, lift it onto the transporter, ship it back to the refurbishment factory, and insert the replacement.

That is fundamentally different from today's industrial plants, where the building and equipment become effectively permanent.

What I would prioritize

If the objective is to cut the **total delivered cost per tonne of chemical**, rather than merely make an interesting transportation system, I would put the first development money into six large-module families:

pre-equipped tunnel modules; pipe/manifold platforms; pump/compressor platforms; modular enclosed conveyors and transfer towers; modular waterfront ship loading/unloading systems; and modular storage systems.

Those six, combined with your **12 × 96 m chemical platforms and double-rail transporter**, would allow much more of the entire industrial site—not merely the chemical reaction equipment—to be manufactured in factories.

I think that is the logical next step in the concept: rather than a **factory-made chemical processing platform**, the objective becomes a **factory-made chemical processing site assembled from standardized giant modules**.

Independent Order-of-Magnitude Check on Program Scale

Prompt to AI: How much would it cost and how long would it take to design (not prototype) all the things referred to in this chat? How much would it cost and how long would it take to build scaled-down 100:1 (or something) prototypes of all the things we talked about in this chat? We would want to test these for months, and look for changes (e.g. corrosion on internal surfaces of pipes). We are not working with nuclear fuel. This would be driven by heat from fossil fuel (e.g. coal in China or natural gas in USA). What is my fuel cost per month to runs prototypes? Answer this question for both USA and China.

Response From AI:

Yes. I think this can be turned into a fairly concrete R&D program.

Budget relationship to the current $1B program

This chapter is an independent bottom-up plausibility check, not an additional budget. The controlled program budget is the 11-work-package $800M technical plan plus the $200M centrally controlled reserve. Country-specific design and prototype estimates below are sensitivity checks on program scale and must not be added to those work-package totals.

My central estimate is that designing the complete system to a serious preliminary/FEED level would cost roughly $180 million in the U.S. or $100 million in China and take about 2–2½ years. Building and operating a useful family of scaled prototypes—not merely tabletop models—would add roughly $500 million in the U.S. or $260 million in China, with about another 2–3 years of overlapping fabrication and testing.

So, a design + scaled-prototype + 12–18-month test program could plausibly fit within a $1 billion program even in the United States, and substantially below $1 billion in China. That excludes construction of a full-scale commercial nuclear/chemical site.

Your business plan already points in this direction: pages 54–56 call for the chemical-processing yard, factory-made platforms, underground tunnels, platform transportation, double-rail equipment, concrete-processing equipment and automated thermal-storage construction. The ideas we subsequently added—material-handling modules, waterfront equipment, storage modules and segmented electrified rail—expand that program.

What I am including

For purposes of estimating "all the things we talked about," I am assuming the R&D program covers:

1. The 31 families of industrial processing platforms we discussed.
2. The common chemical yard: underground tunnels, pipe headers, manifolds, pumps/compressors, conveyors, rail carts, storage and waterfront loading/unloading.
3. The 12 × 24 m double-width railcars, four-car 12 × 96 m transporter, rigid truss, switches, platform insertion/removal and ship/shore transfer.
4. Factory/shipyard methods for constructing the giant platforms and logistics modules.
5. Automated excavation, concrete mixing/pumping, rebar, gantry/crane and tunnel-installation machinery relevant to constructing the yard.
6. Automated construction of large thermal-storage tanks.
7. The segmented 440/690-V and medium-voltage electrified-rail concept, collectors, sectional switchgear, controls and safety systems.
8. Common mechanical/electrical/thermal/process/control standards, digital models, safety analysis, materials qualification, corrosion testing and cost models.

I am **not** including nuclear-fuel systems, reactor physics, reactor licensing or construction of a nuclear reactor.

1. Cost to design everything

By "design" I mean approximately a 30–40% engineering/FEED package: requirements, process flow diagrams, heat/mass balances, CAD layouts, major equipment sizing, structural concepts, controls architecture, safety analyses, interface standards, prototype specifications and cost models. It does not mean complete fabrication drawings for every nut, bolt and weld.

Design work	USA	China
Overall site architecture, digital model and common standards	$15–25M	$8–15M
31 processing-platform designs	$55–80M	$30–45M
Material handling, storage, tunnels, manifolds, conveyors, port modules	$20–35M	$10–20M
Double-rail transport, transporter and ship/shore transfer	$15–25M	$8–15M
Automated excavation/concrete/tunnel construction machines	$20–35M	$10–20M
Automated thermal-storage tank construction	$10–20M	$5–10M

Electrified rail, switchgear, collectors, controls and safety	$5–10M	$3–6M
Materials, HAZOP/safety, corrosion program, reliability and test planning	$10–20M	$5–10M
Total	$150–250M	$80–140M
Central estimate	~$180M	~$100M

These country differences are mostly labor and fabrication-engineering costs. As a reference point, U.S. production/nonsupervisory workers in engineering and drafting services were averaging about $47.77/hour in June 2026 before employer overhead, benefits and engineering-company markup, while engineering managers have a median salary around $168,000/year. China's National Bureau of Statistics reports an average 2025 wage of RMB 155,491/year for professional and technical personnel in larger enterprises. These aren't consulting billing rates, but they illustrate why a labor-intensive design program should cost appreciably less in China.

I would allow 24–30 months in the U.S. and roughly 20–28 months in China if perhaps 200–300 engineers are working in parallel. China should be cheaper; I wouldn't assume it is automatically twice as fast.

If you want production-ready detailed designs, rather than FEED-level designs, I would roughly double the design budget and add perhaps another 12–18 months.

2. I would NOT build everything at 1:100 scale

This is important.

A 1:100 model of a 12 m × 96 m platform is only 0.12 m × 0.96 m.

That's excellent for demonstrating:

- rail movements;
- platform placement;
- site layout;
- switching;
- underground corridors;
- automated sequencing.

It is essentially useless for answering:

- corrosion;
- pipe pressure drop;
- heat transfer;
- pump behavior;

- solids handling;
- thermal cycling;
- seals;
- valves;
- weld durability.

Those phenomena don't simply scale geometrically.

I would therefore build **different scales for different questions**.

Test	Recommended scale
Entire site / choreography	1:100
Railcars, transporter, switches, gantries	~1:10
Conveyors, loading equipment, mechanical modules	~1:5 to 1:10
Process platforms	pilot throughput, perhaps 1:10–1:50 capacity, not geometric scaling
Pipes/valves/materials	Often full-size diameter representative sections
Corrosion loops	Actual material + actual fluid + actual temperature
Thermal-storage construction	Perhaps 1:5–1:10 tank plus full-size welding/inspection trials
Electrified rail	Preferably full voltage, reduced track length and reduced power

For corrosion in particular, I would use removable coupons and spool pieces and pull specimens after perhaps **3, 6 and 12 months**, while continuously monitoring temperature, chemistry and flow. ASTM G31 specifically warns that corrosion results depend on solution composition, temperature, motion, exposure duration and other conditions, and that accelerated testing can be misleading.

So **don't scale the chemistry by 100:1**. Scale the throughput while preserving the relevant physical environment.

3. Cost to build and test the prototypes

I would envision something closer to a small industrial research campus than a collection of models.

Prototype/test component	USA	China
1:100 complete physical yard + digital twin	$2–5M	$1–3M
1:10 double-rail/transporter/gantry/logistics demonstrators	$40–70M	$20–40M
31 functional process-platform pilot systems	$150–300M	$75–160M
Tunnels, manifolds, pumping, conveyors, storage and waterfront	$40–70M	$20–40M

prototypes		
Thermal-storage tank + automated tank-building equipment	$25–50M	$12–30M
Electrified-rail test loop and high-power machinery	$10–25M	$5–15M
Full-temperature flow loops/materials/corrosion facility	$20–40M	$10–20M
Test-site utilities, instrumentation, controls and safety	$30–50M	$15–30M
12–18 months of testing, staff, maintenance and analysis	$40–70M	$20–40M
Total	$360–680M	$180–380M
Central estimate	~$500M	~$260M

These are program-planning estimates rather than contractor quotations.

Schedule

I would expect approximately:

Months 0–12: design progresses; test requirements frozen.

Months 9–24: prototype fabrication begins while detailed design continues.

Months 18–30: test facility and first prototypes commissioned.

Months 24–42: progressively operate the different systems.

Months 30–48: long-duration thermal cycling, corrosion and reliability testing.

Thus, from a standing start:

USA: ~4 years to complete the major scaled-prototype program.

China: perhaps ~3½–4 years.

You don't have to wait until the entire design program finishes before starting prototypes.

4. Total design + prototype program

The central numbers I'd use in the proposal are therefore:

Program	USA	China
Design	~$180M	~$100M
Prototype construction/testing	~$500M	~$260M
Subtotal	~$680M	~$360M

Suggested contingency/reserve	~$120–200M	~$70–120M
Practical program budget	~$800M	~$450M
Calendar duration	~4 years	~3½–4 years

That is an interesting result because it suggests a **$1 billion / 5-year development fund is approximately the right order of magnitude even in the U.S.** In China, the same $1 billion would potentially leave considerable money for progressively larger prototypes or the first near-full-scale demonstrations.

It also agrees directionally with the R&D philosophy in your business plan, which explicitly distinguishes comparatively modest design expenditures from much larger prototype expenditures.

5. Fossil-fuel heat source for testing

Because there is no nuclear fuel, I would make the test facility look like this:

> fossil heat source → intermediate heat loop → prototype

The prototype therefore doesn't know whether its heat originally came from coal, gas or a nuclear reactor.

For PWR-like testing you can provide the appropriate steam/hot-water conditions.

For HTR-PM-like testing, a fossil-fired heater can heat an intermediate fluid or gas to the required higher temperature.

That is much better scientifically than having combustion gases pass directly through the chemical-processing equipment because it preserves the intended nuclear/industrial interface.

6. U.S. prototype fuel cost — natural gas

The latest published EIA U.S. average industrial natural-gas price is **$4.27 per thousand cubic feet for May 2026**. EIA separately forecasts Henry Hub at roughly $3.6/MMBtu in 2026, but an R&D facility pays a delivered industrial price rather than simply the Henry Hub commodity price.

Using:

$4.27/Mcf,
approximately 1.036 MMBtu/Mcf,
and 90% heater efficiency,

I get approximately:

$15.6 per MWh of useful test heat.

A 30-day month contains 720 hours, so:

Average useful heat supplied to prototypes	U.S. gas fuel/month
0.5 MW_{th}	~$5,600
1 MW_{th}	~$11,300
3 MW_{th}	~$34,000
5 MW_{th}	~$56,000
10 MW_{th}	~$113,000

I would provision a prototype campus for roughly 5–10 MWth peak, but I doubt its long-term average needs to exceed roughly 2–3 MWth, because all 31 platform families won't be operating at maximum heat simultaneously.

So, my central U.S. fuel budget would be only:

~$30,000–$40,000/month of natural gas.

Even running at an average 3 MWth continuously for 18 months would consume only around **$0.6 million of gas**.

7. China prototype fuel cost — coal

Chinese coal prices move considerably. A May 2026 market report put 5,500-kcal/kg Qinhuangdao thermal coal around **$122/t**, while prices have moved materially above and below that level. I would therefore use **$90–$120/t** as a reasonable prototype-program sensitivity range rather than pretending today's spot quote will persist throughout a multiyear program.

5,500-kcal/kg coal contains approximately **6.39 MWhth/t**.

Assuming an **85% useful-heater efficiency**, a tonne supplies approximately 5.43 MWh of useful test heat.

That produces:

$17–22/MWh of useful heat at $90–120/t coal.

Therefore:

Average useful heat supplied	China coal fuel/month
0.5 MW_{th}	~$6,000–8,000
1 MW_{th}	~$12,000–16,000
3 MW_{th}	~$36,000–49,000
5 MW_{th}	~$60,000–81,000
10 MW_{th}	~$120,000–162,000

So, I would budget approximately:

> $40,000–$50,000/month of coal for a 3-MW-average Chinese test facility.

Again, that is **fuel only**. Coal ash handling, particulate controls, SO_2/NO_x controls, labor and permitting are separate.

For a relatively small research facility, those coal-related balance-of-plant costs could actually be more bothersome than the cost of coal itself.

8. Fuel is almost irrelevant to the overall R&D budget

This is perhaps the most useful finding.

Take an aggressive case:

> 5 MWth average, 18 months of continuous operation.

Approximate fuel expenditure:

> USA natural gas: ~$1.0M.
>
> China coal: ~$1.1–1.5M.

Compare that with:

> $500M U.S. prototype program or
> $260M Chinese prototype program.

Fuel is around **a few tenths of one percent of the R&D cost**.

The expensive parts are:

> engineers, machines, steel, pipes, pumps, instrumentation, sensors, test stands, controls, repeated modifications, teardown inspections and personnel operating the experiments.

This means you should not scale the experiments down aggressively merely to save fuel. **Run them large enough to obtain trustworthy engineering data.**

What I would actually build

Instead of saying "build everything at 1:100," I would establish one large test campus containing:

> a 1:100 complete-yard demonstrator + a 1:10 transportation/construction yard + several MW of real thermal test loops + roughly 8–12 integrated pilot process skids plus subsystem rigs/test articles for the remaining platform families + a short full-voltage electrified-rail section + a scaled thermal-storage construction area.

Most importantly, I would make the **interfaces full-scale or close to full-scale** wherever the purpose is to learn about connectors, valves, welds, insulation, leaks, corrosion, maintenance or replacement.

The research question isn't whether a pretty miniature factory works. It is whether, after **thousands of hours of hot operation**, you can disconnect Platform #12, inspect the pipe wall, measure corrosion and deposits, move the platform out on the transporter, insert another one, reconnect it, and resume operation economically.

That is the experiment that would make the later full-scale system believable.

Standards Development

Prompt to AI: Companies need to make components for this system that fit together. Therefore, standards are needed that describe how everything described in this chat fits together mechanically, electrically, with heat, and with data communications. This applies to everything in this conversation. What standards are needed? Who might work on these (i.e. standards bodies)? How much might this cost? How long might this take?

Response From AI:

Yes. In this concept, the **interface standards may be almost as important as the machines themselves**. If one company builds a 12 m × 96 m heat-pump platform, another builds the transporter, another the underground tunnel, another the pipe manifold, and another the electrical system, they should be able to connect without redesigning the site.

Budget relationship to Work Package 9

Work Package 9 controls $55M for standards, interoperability, the digital thread, and conformance testing. The U.S./China estimates later in this chapter are external sensitivity ranges for comparable standards work, not extra line items. Additional spending requires an approved reserve release or a documented reallocation under the common change-control rules.

I would **not create new standards for things that already have mature standards**. For example, ASME B31.3 already covers process piping; ASME pressure-vessel rules already exist; China has GB/T 20801 for industrial pressure piping and GB/T 150 for pressure vessels. The new standards should concentrate on the **interfaces unique to your giant modular industrial-site concept**.

My rough estimate is that the complete standards program would involve about **18–22 new interface standards/specifications**, cost roughly **$60–100 million in the U.S. or $35–60 million in China**, incremental to the prototype program we discussed, and require roughly **18 months to get a usable first-generation interface specification, 3 years for mature national standards, and 4–5 years for broad international ISO/IEC standards**.

I would organize the standards into three layers

- **Layer 1 — existing engineering codes.** These specify how a pipe, pressure vessel, motor, crane, weld, electrical enclosure, etc. is safely designed.
- **Layer 2 — new common interfaces.** These specify exactly where the pipe connection is, where the support pad is, what voltage is supplied, what data messages mean, how the transporter aligns, and so forth.

- **Layer 3 — equipment profiles.** A heat-pump platform, ammonia platform, tank-building machine, conveyor module, etc. would have a profile saying which standardized interfaces it uses.

That approach prevents the standards effort from trying to reinvent all of industrial engineering.

The new standards I think you need

Proposed standard	What it would define	Important existing standards to build upon
1. Overall System Architecture	Vocabulary; module classes; system boundaries; common coordinate system; hierarchy of interfaces	ISO industrial automation/system-integration work
2. Giant Module Dimensional Standard	12 × 24, 12 × 48, 12 × 72, 12 × 96 m envelopes; allowable overhangs; heights; weight classes; CG limits	New standard
3. Module Structural Interface	Standard support feet, hard points, lifting points, jacking points, tiedowns, structural load paths	AISC/structural standards; ASME BTH/B30 for lifting
4. Platform-Bay Docking Standard	Exact coordinates of supports; alignment pins; mechanical locks; tolerances; automatic docking	New standard
5. Double-Rail Transportation Standard	Four-rail geometry, gauge, rail loads, curves, switches, braking, coupling, transporter geometry	ISO/TC 269 railway standards
6. Platform Installation/Removal Standard	How a transporter brings a module in, transfers load to permanent supports, disconnects and removes it	New standard
7. Giant Lifting/Gantry Interface	Standard lift points, spreader/truss interfaces, load monitoring and handling procedures	ASME B30/BTH
8. Process-Pipe Interface	Standard positions, diameters/classes, flange types, double isolation, purge/drain, leak detection	ASME B31.3/B16; China GB/T 20801
9. Process-Material Manifold Standard	How a platform identifies and connects feedstocks, products, gases, slurries and return streams	New standard using existing piping codes
10. Nuclear/Industrial Heat Interface	Standard heat grades, temperature/pressure ranges, supply/return, isolation loop, metering, thermal expansion	ASME B31.1/B31.3 and heat-exchanger codes
11. Site Utility Interface	Cooling water, demin water, firewater, nitrogen, compressed air, vacuum, drains, vent collection	Existing industrial codes + new interface geometry
12. Platform Electrical Power Standard	LV/MV voltage classes, connectors, grounding, fault levels, power quality, metering	IEC machinery and switchgear standards
13. Electrified-Rail Standard	Enclosed conductor geometry, collector, segment switching, grounding, illumination, interlocks, emergency isolation	IEC TC 9, IEC 62271, IEC 60529, IEC 62061
14. Controls/Data	Physical network, addressing, time synchronization,	OPC UA/IEC 62541, IEC

Proposed standard	What it would define	Important existing standards to build upon
Communications Standard	commands, alarms, equipment discovery	61158/61784
15. Common Equipment Data Model	Meaning of tags such as temperature, flow, availability, operating mode, maintenance state	OPC UA Companion Specification + ISO 15926
16. Functional-Safety Standard	E-stops, safe states, interlocks, machine authority, personnel exclusion zones	ISO 12100, IEC 61508/62061
17. Cybersecurity Standard	Authentication, certificates, software updates, network zoning, remote access	IEC 62443
18. Tunnel and Riser Standard	Tunnel dimensions, dedicated utility zones, drainage, ventilation, cable routing, material rail/conveyor interface	Civil/fire codes + new interface
19. Bulk-Solids Interface	Conveyor width/capacity, hoppers, chute locations, dust containment, weighing and sampling	Existing conveyor standards + new geometry
20. Waterfront/Ship Interface	Liquid/gas manifolds, loading arms, ESD link, solid unloaders, conveyors, ship-to-platform transfer	ISO/TC 8 and OCIMF
21. Storage/Thermal-Storage Interface	Tank nozzles, pumps, level/temperature instrumentation, foundations and modular construction equipment	ASME TES-1, API tank standards
22. Digital Thread & Conformance Standard	Unique module IDs, CAD data, materials, weld/NDE history, software versions, test certificates and acceptance tests	ISO 15926/ISO 10303-type lifecycle data concepts

Several of these could eventually become parts of one standards family rather than 22 separate books.

The most important standards are probably #2, #4, #8, #10, #12, #13 and #14

Those are what make equipment from different companies physically interchangeable.

For example, imagine three companies:

- **Company A:** builds a 12 × 96 m evaporation platform.
- **Company B:** builds a 12 × 96 m heat-pump platform.
- **Company C:** builds a 12 × 96 m ammonia platform.

They should all place their mechanical feet, electrical connections, communications connections, heat connections and common utility connections at **exactly the same standardized locations**.

A simplified end of every platform might effectively look like:

STANDARD INTERFACE END

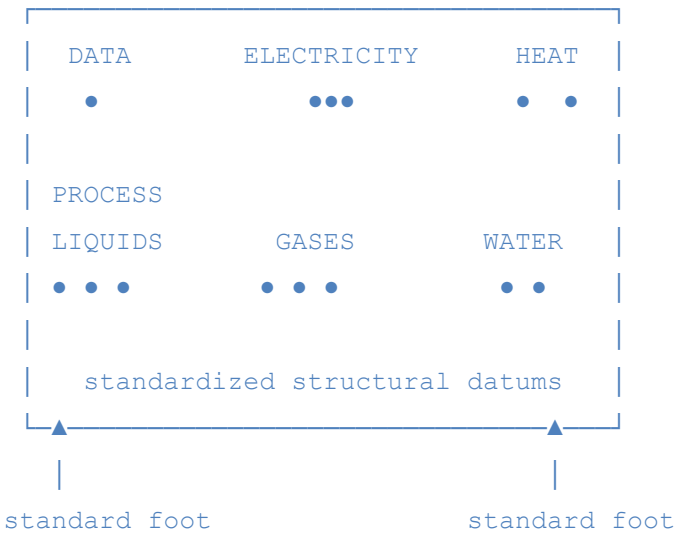

That is analogous to what a USB connector or shipping container does, only at enormous industrial scale.

Heat probably needs its own particularly careful standard

I wouldn't simply define "heat connection."

I would establish several **heat-interface classes**. The exact numbers should come from the engineering program, but conceptually they might be:

Heat class	Approximate range	Likely source/use
H1	~60–150°C	heat pumps, waste heat, hot water
H2	~150–300°C	PWR steam/process heat
H3	~300–550°C	HTR-PM process heat
H4	>550°C	electrically topped/high-temperature systems

The standard would specify much more than temperature:

> medium → pressure → maximum/minimum flow → supply temperature → return temperature → connection position → allowable pressure drop → isolation → leak detection → thermal metering → emergency shutdown → purity requirements.

A company could then advertise:

Platform XYZ requires: H2-200 MW thermal + E1-20 MW electrical + CW2 cooling-water interface.

That would make plant design almost catalog-like.

The electrified rail deserves a new standard of its own

There are existing standards covering railway electrical systems, machinery, medium-voltage switchgear and electrical enclosures, but the **segmented, normally-dead, enclosed conductor rail for multi-megawatt industrial construction machines** is sufficiently different that I would create a dedicated interface standard. IEC TC 9 covers electrical equipment and fixed installations for railways, while ISO/TC 269 covers the non-electrical railway system and infrastructure interfaces.

The new specification should define things such as:

> 6.6-kV versus lower-voltage classes; 3-phase conductor + protective earth; conductor spacing; downward-facing access slot; collector geometry; collector contact force; permissible speed; electrical-segment length; automatic segment energization; machine-position verification; ground fault detection; arc-fault protection; emergency grounding; lockout/tagout; illumination of live sections; redundant proof that a segment is dead before personnel access.

For the medium-voltage switching hardware, IEC 62271 already covers AC switchgear above 1 kV and up to 52 kV, so I would reuse it rather than invent new circuit breakers. IEC 60529 provides enclosure/IP protection requirements, and IEC 62061 provides functional-safety requirements for machinery controls.

China already uses the same general approach of incorporating international electrical standards: **GB/T 5226.1-2019 is an identical adoption of IEC 60204-1:2016** for electrical equipment of machinery.

That is a good reason to aim for **one international conductor-rail interface**, rather than incompatible Chinese and American versions.

Data communication should probably use OPC UA

I would strongly prefer **not** to invent another industrial communications protocol.

IEC already recognizes the IEC 62541 OPC UA family for industrial communications, and the OPC Foundation specifically supports **Companion Specifications** that add an industry-specific semantic information model on top of OPC UA.

I would therefore develop something like an:

"OPC UA Companion Specification for Modular Industrial Processing Sites."

Every compliant platform would expose the same objects:

```
Module
 ├─ Identity
 ├─ MechanicalStatus
 ├─ ElectricalInterface
 ├─ HeatInterface
 ├─ MaterialInputs
 ├─ MaterialOutputs
 ├─ SafetyState
 ├─ OperatingMode
 ├─ Alarms
 ├─ Maintenance
 └─ RemainingUsefulLife
```

The central site computer could discover a newly installed platform automatically.

ISO 15926 is also relevant because it is specifically intended for **integration of lifecycle data for process plants**, while ISA-95/IEC 62264 addresses integration of manufacturing/logistics systems with plant control.

So, I would use:

> OPC UA for live communications

and

> ISO 15926-type definitions for long-term engineering/lifecycle data.

Cybersecurity must be in version 1.0

Because these machines could move thousands of tonnes, switch megawatts of power, or open chemical valves, cybersecurity cannot be added later.

IEC 62443 already provides industrial automation cybersecurity requirements, including secure product-development requirements and security requirements for industrial components and systems.

The new standard should therefore specify how IEC 62443 is applied to the modular yard rather than inventing a new cybersecurity framework.

For instance, a replacement platform arriving from the factory shouldn't be allowed onto the operating network merely because someone plugged in its fiber connector.

It should present a cryptographic machine identity, authenticate, load its certified data model, and be granted only the commands appropriate to that equipment.

Thermal storage already has a useful standards home

This part is unusually convenient.

ASME has a dedicated **Thermal Energy Storage Standards Committee**, and **ASME TES-1** already establishes lifecycle requirements for molten-salt thermal storage, including storage vessels, pumps, valves, piping, heat exchangers and controls.

Therefore, your automated tank-building idea probably does **not** need a new standard for basic molten-salt storage safety.

It needs a new standard for:

> factory-made tank panels + automated tank construction + robotic welding/NDE + standardized tank-service modules + connections to the site heat network.

ASME would be a particularly logical home for that work.

Ship loading/unloading also has a substantial existing foundation

ISO/TC 8 already has committees specifically covering **ship piping/machinery, fluid-transfer connections, ship-to-shore interfaces, ports and terminals**.

OCIMF also maintains detailed guidance for marine loading arms, including operating envelopes, products transferred, manifold spacing and jetty/piping layouts.

Therefore, the new standard should concentrate on the **yard-side modular interface**, rather than redesigning the ship's loading arm.

For example:

> OCIMF-compliant loading arm → standardized Waterfront Module Interface WM-1 → standardized underground material spine → standardized Platform Material Interface PM-1.

That allows the new system to plug into established marine practice.

Who should develop the standards?

I would not give the whole job to one standards organization. It spans too many specialties.

Area	Organizations I would involve
Overall architecture	ISO, ANSI, SAC; industrial consortium
Giant platforms / structures	ASME, AISC, ISO
Railcars / track / transporter	ISO/TC 269
Electrified rail	IEC/TC 9, IEC electrical committees
Process piping	ASME B31, ISO/TC 67; Chinese pressure-piping committees
Pressure vessels	ASME BPVC, Chinese GB/T 150 committees
Thermal storage	ASME TES Committee
Machinery safety	ISO/TC 199, IEC
Electrical power	IEC, IEEE; NFPA for U.S. installation/safety implementation
Industrial automation	IEC/TC 65, ISA
Data interoperability	OPC Foundation, ISO/TC 184/SC 4
Cybersecurity	IEC/ISA 62443 groups
Ships/port interface	ISO/TC 8, OCIMF
Tanks/liquid terminals	API, ASME
China implementation	SAC/SAMR and relevant Chinese national technical committees
U.S. implementation	ANSI-coordinated standards developers
European implementation	CEN/CENELEC; ETSI where relevant; national standards bodies and applicable European nuclear, pressure-equipment, machinery, hazardous-area, cyber, marine, and industrial technical committees

ANSI itself does not normally write the technical standard; it coordinates the U.S. voluntary standards system and accredits standards-developing organizations. In China, the national standards platform is administered by SAMR/SAC and routes standards through national technical committees.

I would therefore create a **project standards consortium** first, with formal liaisons to those organizations.

Don't wait three years for ISO

This is important.

ISO says development of a conventional International Standard usually takes about **three years**, and ISO projects are normally placed on 18-, 24- or 36-month development tracks after formal initiation.

You cannot tell 50 companies:

Don't design your equipment until ISO finishes five years from now.

Instead, I would use this sequence:

Time	Deliverable
0–3 months	Standards steering committee; define system architecture
3–6 months	Freeze module dimensions, coordinate system and basic interface philosophy
6–9 months	Draft mechanical, heat, electrical, rail and communications interfaces
~12 months	Interface Specification v0.9 used by prototype suppliers
~18 months	Open Interface Standard v1.0 — stable enough for commercial design
18–30 months	Interoperability testing with multiple companies
~30–36 months	Revised v2.0 / national consensus standards
~36–60 months	ISO/IEC international standards publication

ISO permits faster preliminary deliverables as well as full International Standards, but the formal process still involves proposal, drafting, enquiry, voting and publication stages.

So, **18 months to a usable industry specification** is realistic if it is funded aggressively.

Formal international consensus takes longer.

How much would this cost?

Standards organizations generally don't charge "$50 million to write a standard." Much of normal standards work is contributed by participating companies and experts.

Your situation is different because you would want standards to be developed **before an industry exists**. Someone therefore needs to pay engineers to establish the interfaces, make reference CAD models, build test tools and prove interoperability.

My planning estimate would be:

Standards-development activity	USA	China
Master architecture + systems engineering	$5–8M	$3–5M

Standards-development activity	USA	China
Mechanical/module/rail interface standards	$8–15M	$4–8M
Piping/material/heat standards	$6–10M	$3–6M
Electrical + electrified-rail standards	$6–12M	$3–7M
Controls/data/cyber standards	$5–10M	$3–6M
Port/storage/material-handling standards	$5–8M	$3–5M
Digital models/reference drawings/schemas	$5–10M	$3–6M
Interoperability + conformance testing*	$10–20M	$7–12M
Standards secretariat, legal/IP, publication, international coordination	$5–10M	$3–5M
Total	**~$55–103M**	**~$32–60M**

*This assumes the physical prototype program we discussed previously already exists. I am **not** charging the standards program again for hundreds of millions of dollars of prototype hardware.

I would round this to:

USA: ~$75 million central estimate

China: ~$45 million central estimate

for a serious standards-and-interoperability program.

That is only around **7–10% of a $1 billion equipment-development program**, which strikes me as a very sensible allocation.

You could spend much less—but I wouldn't

If the objective were merely to write documents, perhaps:

USA: ~$20–35M

China: ~$10–20M

could produce specifications.

But I wouldn't trust them.

The critical work is making Company A's platform physically connect to Company B's tunnel, Company C's rail transporter, Company D's power system and Company E's control system **without engineers standing around with grinders, adapters, custom cables and new software**.

That requires real interoperability testing.

I would build a standards interoperability laboratory

This should probably be part of the prototype campus we discussed.

It would contain a representative:

> platform bay + double rail + standard support feet + pipe manifold + H1/H2/H3 heat connections + LV/MV power + electrified rail segment + fiber/data connection + material conveyor interface.

Suppliers bring equipment in.

The test becomes:

> Can we drop Company X's module into Bay #4 and connect it without modifying either the module or the site?

If yes, it passes.

If not, either the manufacturer isn't compliant or the standard is deficient.

The OPC Foundation already uses certification/compliance testing as part of its interoperability ecosystem, which is a useful model for the data side of this system.

Eventually you could have certification labels such as:

> Mechanical M2 compliant
> Heat H2/H3 compliant
> Electrical E2 compliant
> Data D1 compliant
> Material P3 compliant
> Transport R4 compliant

A company could then sell a machine without knowing exactly which Chinese nuclear-industrial site it will ultimately go to.

There is an even bigger economic implication

The purpose of these standards isn't merely convenience.

It is to create **competition at the component level**.

Without standards:

> Company A designs the reactor-side interface, platform, piping, controls and transport equipment as one proprietary system.

You essentially have one giant EPC supplier.

With standards:

> 10 companies can make heat pumps.
> 6 companies can make evaporator platforms.
> 4 companies can make giant transporters.
> 20 companies can make pumps.
> 5 companies can make thermal-storage tanks.
> 8 shipyards can fabricate 12 × 96 m platforms.

All can bid against one another because the interfaces are fixed.

That is the same underlying economic phenomenon that made **shipping containers, USB, Ethernet and standardized electrical systems** so powerful: innovation can occur **inside the box while the interface stays stable**.

For this project, I would therefore make the standards an explicit major work package—probably **$50–75 million in China or roughly $75–100 million in the U.S., beginning on day one and running in parallel with machine and platform design**. The first interface specifications should be frozen sufficiently by about **month 12–18** for prototype manufacturers to build against them, with national consensus standards around year 3 and international ISO/IEC standardization following around years 4–5.

Direct Air Capture (DAC) Using Nuclear Power

Prompt to AI: If Automated China PWR or Automated China HTR-PM supplies the electricity and heat for direct air capture, what is the resulting cost per tonne of CO_2 permanently removed? Use the nuclear energy-cost assumptions already adopted in this document and distinguish energy cost from total DAC cost.

Response From AI:

Using the Automated China PWR and HTR-PM planning assumptions in this document, nuclear-powered DAC becomes substantially more attractive than today's first-of-a-kind DAC economics. The strongest configuration is likely a low-temperature solid-sorbent DAC process that uses electricity for fans, vacuum/compression and controls and roughly 80–120°C heat for sorbent regeneration.

The nuclear cost basis is the same one used elsewhere in this proposal: Automated China PWR at approximately \$31/MWh_e and \$10/MWh_{th}, and Automated China HTR-PM at approximately \$37/MWh_e and \$15/MWh_{th}. These remain scenario/R&D target values, not demonstrated commercial prices.

Recommended mature DAC planning values

Nuclear + DAC configuration	Nuclear energy cost inside DAC	Estimated DAC capture cost	CO_2 transport + permanent storage	Total durable CO_2 removal
Automated China PWR + solid-sorbent DAC	~\$23–50/tCO_2	~\$90–130/tCO_2	~\$10–20/tCO_2	~\$100–150/tCO_2
Automated China HTR-PM + solid-sorbent DAC	~\$29–70/tCO_2	~\$100–140/tCO_2	~\$10–20/tCO_2	~\$110–160/tCO_2

For program-level planning, I would use central values of approximately \$120/tCO_2 permanently removed for Automated China PWR and \$130/tCO_2 for Automated China HTR-PM, with roughly \$80–180/tCO_2 as a broad sensitivity range. These are mature, nth-of-a-kind targets rather than current commercial DAC prices.

Energy calculation: advanced low-energy DAC case

Assume an advanced solid-sorbent DAC system eventually requires about 2 GJ/tCO_2 of electricity and 2 GJ/tCO_2 of low-temperature regeneration heat. Because 1 MWh = 3.6 GJ, each term is 0.556 MWh/tCO_2.

Energy term	Automated China PWR	Automated China HTR-PM
Electricity: 0.556 MWh/tCO_2	0.556 × \$31 = \$17.2/tCO_2	0.556 × \$37 = \$20.6/tCO_2
Heat: 0.556 MWh_{th}/tCO_2	0.556 × \$10 = \$5.6/tCO_2	0.556 × \$15 = \$8.3/tCO_2
Total nuclear energy	~\$23/tCO_2	~\$29/tCO_2

The result is important: with very inexpensive nuclear electricity and heat, energy itself need not dominate DAC cost. The remaining cost is primarily air contactors, sorbent inventory/replacement, vacuum equipment, fans, heat exchangers, structures, compression equipment, maintenance, labor and capital recovery.

Sensitivity: higher regeneration-heat requirement

DAC heat requirements remain uncertain. As a deliberately less favorable sensitivity case, use about 2 GJ/tCO_2 of electricity but 11.9 GJ/tCO_2 of regeneration heat. The electricity term remains 0.556 MWh/tCO_2, while 11.9 GJ of heat equals about 3.31 MWh_{th}/tCO_2.

Energy term	Automated China PWR	Automated China HTR-PM
Electricity	~\$17/tCO_2	~\$21/tCO_2
Regeneration heat	3.31 × \$10 = ~\$33/tCO_2	3.31 × \$15 = ~\$50/tCO_2
Total nuclear energy	~\$50/tCO_2	~\$70/tCO_2

This sensitivity illustrates why DAC R&D should focus heavily on lowering regeneration energy and sorbent cost. Moving from a high-heat cycle toward a few GJ/tCO_2 can reduce the nuclear energy contribution by tens of dollars per tonne.

Why Automated China PWR is the preferred DAC reactor

HTR-PM supplies much higher-temperature steam, but a low-temperature solid-sorbent DAC plant does not need 500°C-class heat. Under this document's planning assumptions, PWR electricity is cheaper (\$31 versus \$37/MWh_e) and PWR heat is cheaper (\$10 versus \$15/MWh_{th}), while PWR steam is already hot enough to supply an isolated ~80–120°C DAC regeneration loop. Therefore, for a stand-alone DAC complex, Automated China PWR is the lower-cost nuclear source.

HTR-PM becomes more interesting when heat is cascaded. A 300–500°C industrial platform could use the high-quality HTR-PM heat first; its lower-temperature reject/waste heat could then feed a DAC regeneration loop. In that integrated configuration, DAC may use heat that would otherwise be rejected, improving system economics beyond the stand-alone comparison.

Scale implication

Permanent removal rate	At $120/tCO$_2$ (PWR central case)	At $130/tCO$_2$ (HTR-PM central case)
1 GtCO$_2$/year	$120 billion/year	$130 billion/year
10 GtCO$_2$/year	$1.2 trillion/year	$1.3 trillion/year

Even at approximately $100–150/tCO$_2$, multi-gigaton DAC is a very large industrial undertaking. The potential value of the nuclear/DAC combination is therefore not that removal becomes cheap in an absolute sense, but that inexpensive 24/7 heat and electricity may make durable, precisely measurable removal substantially less costly than first-of-a-kind DAC today.

DAC platform architecture

The platform chapter supports two repeatable DAC modules: (26) an air-contactor + solid-sorbent regeneration platform and (27) a CO$_2$ dehydration/compression/storage-interface platform. Many Platform 26 trains can operate in parallel and feed a smaller number of common Platform 27 trains. This mirrors the broader program philosophy: standardized factory-built equipment, repeatable interfaces, shared utilities and measured cost/reliability data.

Important qualifications

- The Automated China PWR and HTR-PM electricity/heat prices are scenario estimates from the nuclear-construction proposal and have not been commercially demonstrated.
- The approximately $100–160/tCO$_2$ total-cost ranges are mature planning targets. They require substantial reductions in DAC capital cost, sorbent cost, energy consumption and maintenance relative to many current systems.
- Net removal must subtract lifecycle emissions from construction, replacement materials, transportation and energy. CO$_2$ must then be stored durably, such as in qualified geologic formations.
- A DAC development program should separately measure electricity use, heat use, sorbent degradation, water use, air-contactor pressure drop, capture efficiency, CO$_2$ purity, plant availability, compression energy and verified net tonnes permanently stored.

Links to references used for the DAC planning calculation

- https://www.osti.gov/servlets/purl/2520078
- https://www.energy.gov/sites/default/files/2024-11/Carbon%20Negative%20Shot_Technological%20Innovation%20Opportunities%20for%20CO2%20Removal_November2024.pdf
- https://pubs.acs.org/doi/10.1021/acs.est.5c14628

- https://www.sciencedirect.com/science/article/pii/S175058362500088X

Automated Construction of Thermal Storage Tanks

Prompt to AI: In theory, thermal storage tanks could be made using super-sized automated equipment (e.g. on 12m x 24m railcars). How might this be done? Please see page 56 of "Decarbonization_Lab_Biz_Plan_v8h.docx" for comments on this.

Response From AI:

Page 56 is brief, but the idea is clear: use thermal storage to absorb heat when the reactor has more heat than immediately needed, then recover that heat later; and reduce storage cost by developing **custom machines that automatically construct very large tanks from factory-made subassemblies**.

I think this could fit your double-rail concept very well. The key is **not to transport the completed tank**. A truly useful thermal-storage tank might be 40–100 meters in diameter. Instead, transport very large tank components and the construction machines on 12 m × 24 m double-width railcars, then have the machines assemble the tank almost like a giant automated kit.

A good basic architecture

For high-temperature sensible-heat storage, I would probably start with either:

- **two large tanks** — one "hot" and one "cold"; or
- **one thermocline tank** with hot liquid at the top and cooler liquid at the bottom.

The exact storage fluid would depend on the reactor and temperature. The construction technology can be designed largely independently of whether the eventual fluid is a molten salt, thermal oil, or another heat-transfer fluid.

A pair of tanks might look conceptually like this:

> HTR-PM/PWR → heat exchanger → hot tank → industrial processes/turbine → cold tank → reactor heat exchanger

The interesting R&D problem is then: **How cheaply can we build tanks of tens of thousands of cubic meters repeatedly?**

I would use a "build from the top downward" tank-jacking system

There is a particularly attractive way to automate this.

Instead of erecting the bottom shell first and having workers climb higher and higher as the tank grows, build the **roof and top wall ring at ground level**.

Then:

- Construct the tank floor.
- Build the roof and top cylindrical shell ring at ground level.
- Put perhaps 30–100 synchronized hydraulic/electric jacks around the circumference.
- Lift the entire roof/top ring several meters.
- Bring in the next ring of factory-made curved wall panels underneath it.
- Automatically align and weld that ring.
- Raise the entire structure again.
- Repeat until the full tank height has been created.

That means nearly all difficult joining occurs **close to ground level** rather than 20 or 30 meters in the air.

I think this is particularly compatible with your automation philosophy.

What comes in on the 12 m × 24 m railcars?

Instead of ordinary truck-sized construction material, a double-width railcar could carry things like:

Railcar load	What factory does beforehand
Large curved wall panels	Cut, roll to exact curvature, edge-machine, inspect
Tank-floor sectors	Cut and weld large subassemblies
Roof sectors	Fabricate large radial roof sections
Insulation modules	Factory-build insulation + outer weather cladding
Pipe/manifold module	Pumps, valves, heat exchangers, instruments preassembled
Automated welding machine	Complete robotic fabrication cell
Tank-jacking machine	Large synchronized jack modules
Foundation/rebar modules	Factory-made reinforcement/forms
Concrete-processing car	Aggregate/cement/water bins + automated mixer/pump
NDE/inspection machine	Automated ultrasonic/radiographic/weld inspection

This follows the page-56 proposal very closely: **large factory-made subassemblies joined on site by custom machines**, using the same double-rail infrastructure being proposed for the rest of the chemical-processing yard.

The tank-construction "train"

I would probably develop several specialized machines, each mounted on one or more 12 × 24 m railcars.

1. Foundation Construction Car

This prepares the tank base.

It could carry:

- laser surveying;
- automated excavation/grading tools;
- rebar placement;
- concrete mixing/pumping equipment;
- automated screeding;
- embedded anchors;
- foundation insulation installation.

For a high-temperature tank, the foundation is important because you don't want heat damaging the concrete or migrating excessively into the soil.

The machine could make a standardized circular foundation perhaps 50–80 m across.

2. Floor Assembly Machine

The factory ships **large floor plates or pie-shaped floor sectors**.

A railcar-mounted gantry lifts them off the transporter and lays them onto the foundation.

Then automated weld heads travel along every seam.

Rather than hundreds of workers manually fitting plates:

> place → align → clamp → weld → inspect

would be machine-controlled.

3. Giant Tank-Jacking System

Around the perimeter, standardized jack modules would be placed perhaps every few meters.

They operate synchronously.

The sequence becomes:

> lift 3 m → stop → insert new shell ring → weld → inspect → lift another 3 m.

The important point is that **the tank itself moves vertically while the welding machines stay at approximately ground level**.

That greatly simplifies automation.

4. Shell-Panel Placement Machine

This might be one of the most visually impressive machines.

Imagine two 12 × 24 m railcars, one on either side of the construction area, supporting a huge gantry.

A trolley travels across the gantry.

The trolley picks up a large curved steel panel from a delivery railcar and positions it against the tank.

Robotic clamps then:

- hold the panel;
- measure alignment with lasers;
- adjust millimeters;
- tack weld it;
- release the gantry.

The next panel comes in.

Once a complete circular ring exists, an automated circumferential welder completes the vertical and horizontal seams.

5. Automated Welding Machine

This deserves its own development program because tank construction involves **an enormous amount of repetitive welding**.

The welding machine could magnetically clamp or mechanically attach to the tank wall and travel around it.

It could perform:

- seam preparation;

- preheating where necessary;
- welding;
- controlled cooling;
- grinding;
- inspection.

Multiple welding heads could work simultaneously.

For example, instead of one worker welding one seam:

> 20 robotic heads could weld 20 seams at the same time.

That is where a serious construction-cost reduction could emerge.

6. Automated inspection immediately behind the welding machine

I would make welding and inspection one integrated operation.

Something like:

> Welder → ultrasonic scanner → machine vision → dimensional scanner

follows the weld.

Software records every centimeter of every seam.

If a defect is found, its coordinates are automatically tagged and a repair robot comes back.

That could actually produce **higher quality** than traditional manual tank construction while lowering labor cost.

7. Roof Construction

The roof can be factory-made in very large wedge-shaped sections.

A central temporary column or erection tower could support the roof during assembly.

Alternatively, because we're using the jacking method, I would assemble nearly the **entire roof at ground level before raising the tank**.

That is a major advantage.

Workers and robots aren't installing roof pieces 25 meters above the ground.

8. Insulation Installation Machine

After the steel tank passes pressure/leak testing, another automated machine works around its exterior.

Large factory-made insulation panels arrive by double rail.

A robotic gantry installs:

> thermal insulation → vapor/weather barrier → metal outer cladding.

Again, work is repetitive and geometrically simple—exactly what automation likes.

9. Factory-built piping and pump platform

I wouldn't field-build the complicated equipment around the tank.

Beside each tank there could be a **12 × 24 m or 12 × 48 m tank-service platform**, factory completed with:

- main circulation pumps;
- valves;
- drain system;
- fill system;
- emergency drain;
- instrumentation;
- level sensors;
- heat-tracing equipment;
- electrical equipment;
- heat exchangers.

The tank then has perhaps only 5–10 large standardized connection points.

This follows the same plug-and-play philosophy as the chemical platforms.

The rail layout could be part of the construction machine

I envision the construction site roughly like this from above:

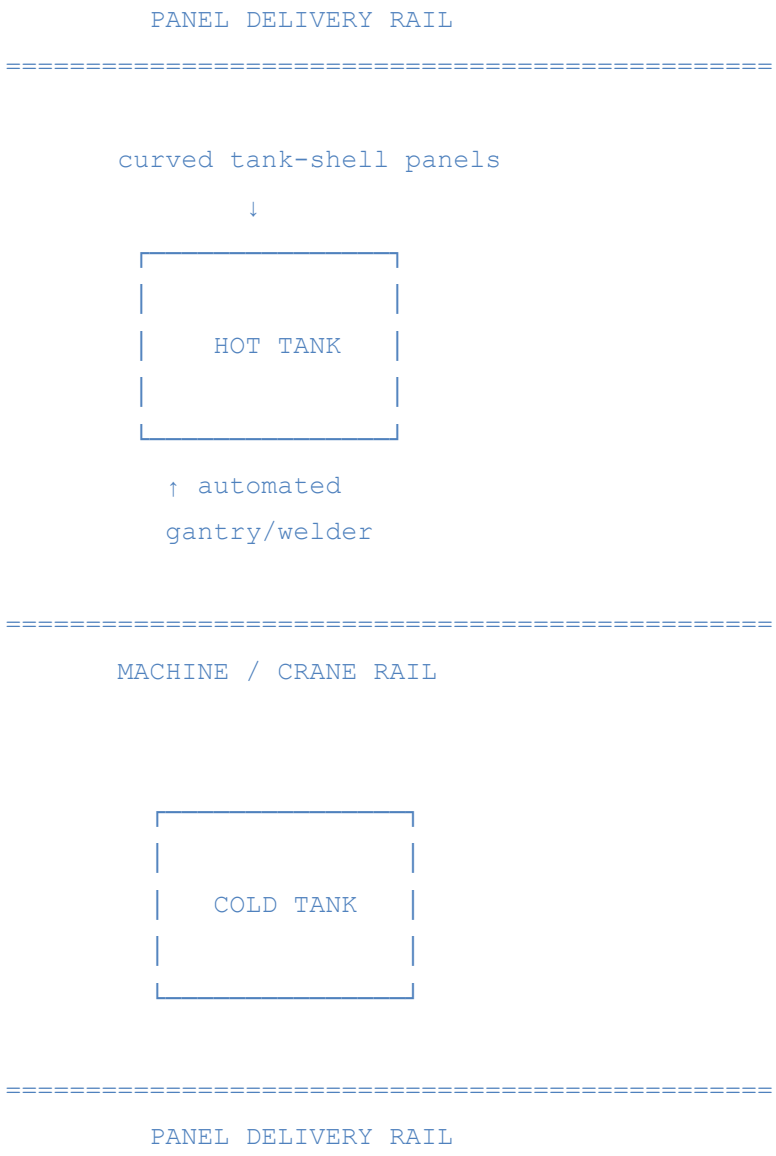

The same heavy-double-rail infrastructure used to bring in chemical-processing platforms can therefore support construction and maintenance of thermal storage.

An alternative: manufacture the steel wall continuously on site

There is another approach worth testing.

Instead of shipping curved shell panels, ship **large rolls or stacks of flat steel plate**.

A giant railcar-mounted machine would contain:

steel feed → edge preparation → rolling machine → curvature control → positioning → welding

and effectively manufacture the cylindrical wall continuously at the tank site.

Think of it as a gigantic tank-printing machine.

It may or may not beat factory-produced curved panels. I would develop both concepts on paper and cost them.

My expectation would be:

factory-curved panel + automated field assembly wins initially because it maximizes controlled factory fabrication.

But the continuous construction machine could eventually become cheaper at very high tank-production volumes.

Concrete can also be automated

The other possibility is that the containment/structural portion of the tank is mostly concrete, with a separate metal liner.

That fits another item already on page 56: the business plan proposes **double-rail concrete-processing equipment**, including railcars carrying water, aggregate and dry cement and machines that mix and place concrete.

One could therefore have:

```
Railcar #1 — cement
Railcar #2 — aggregate
Railcar #3 — water/additives
Railcar #4 — mixer/pump
Railcar #5 — automated rebar/form machine
Railcar #6 — concrete-placement gantry
```

For very large storage systems, a cheap concrete outer structure plus a factory-made metallic liner might deserve comparison against an all-steel tank.

Very large size could improve economics

Suppose, just illustratively, that a tank was:

```
60 m diameter × 20 m liquid depth
```

Volume is about $\pi \times (30\ \text{m})^2 \times 20\ \text{m} \approx 56{,}550\ \text{m}^3$.

A storage liquid with density around 1.7 t/m³ would mean roughly **96,000 tonnes of storage medium**.

A site might therefore have something like:

```
Hot Tank A / Cold Tank A
Hot Tank B / Cold Tank B
```

and add additional pairs as reactor and chemical-processing capacity grows.

Thermal storage becomes more useful with your chemical yard

Page 56 discusses using storage when electricity demand is low and later generating electricity when demand rises. It also explicitly notes that storage can absorb otherwise-unused heat.

With your expanded chemical-yard concept, I would broaden the control logic further.

The storage system could balance **three demands simultaneously**:

```
electricity demand
chemical process-heat demand
reactor output
```

For example:

```
                          → Electricity turbine
                         /
Reactor → heat header → Chemical platforms
                         \
                          → Thermal storage
```

At 2 AM:

1. grid electricity demand falls;
2. chemical plants continue;
3. surplus reactor heat charges storage.

At 4 PM:

- grid electricity price rises;
- stored heat drives additional turbine generation.

If a chemical platform unexpectedly shuts down:

- instead of reducing reactor output or dumping heat,
- divert its thermal load temporarily to storage.

If another chemical platform needs a burst of heat:

- reactor + thermal storage supply it together.

So, storage can also permit **smaller reactor-side equipment** relative to peak process demand.

I would develop perhaps seven specialized machines

If this became an R&D package inside your fund, my shortlist would be:

- Automated tank-foundation construction machine
- Tank-floor panel placement/welding machine
- Synchronized tank-jacking system
- Shell-panel positioning gantry
- Automated welding + NDE machine
- Roof assembly machine
- Automated insulation/cladding machine

All would share the **12 × 24 m double-rail transportation standard**, and most could remain at the nuclear industrial complex and move from tank to tank.

That last point is important economically. **You don't buy a giant automated tank-building machine for each tank.** You buy perhaps one set of machines for a site, build Tank #1, roll the machines 100 meters down the tracks, build Tank #2, then Tank #3, and so forth.

If China ultimately constructed hundreds of nuclear/chemical super-sites, the machines themselves could become standardized manufactured products.

That, I think, is the strongest version of the page-56 idea: **don't optimize the design of one thermal-storage tank; design a reusable automated tank factory that happens to move around the site on double rail.**

Electrified Rail to Assist in Site Construction

Prompt to AI: In theory, double-rail on site could include a parallel high voltage rail (e.g. 440 to 7040 VAC, 100 to 1600 Amps, 0.1 to 10 MW). This could be near the ground, next to the rail that supports railcars. Also, it could be surrounded by insulation, except a groove on the bottom. This is illustrated in file “pic3_with_electrical_rail.png” where the electrified rail is red, the insulation is green and a stem that connects the railcar to the electrified rail is shown in blue. Also, there are many segments of rail on the site, and most of them would not need to be electrified at any one time. In other words, mechanical relays could attach power to rail segments as needed. Also, lights on top of electrified rails could illuminate when electrified, to alert workers to danger. This system would enable grid power to support automated equipment that builds these site (e.g. excavation, concrete mixing/pumping, movement of material/machines). What do you think about this?

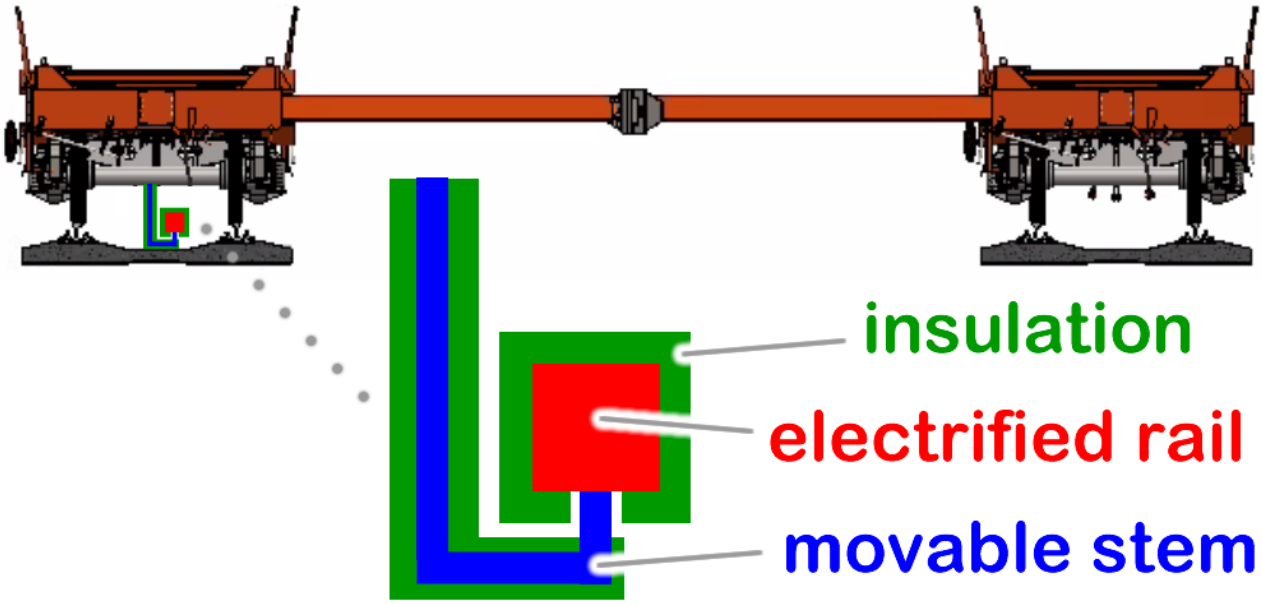

Figure 4. Electrified rail concept.

Response From AI:

I think the **electrified double-rail idea is technically plausible and potentially quite valuable** for your construction concept. I would modify the details somewhat for safety and for multi-megawatt equipment.

In your sketch, the key concept is good: the normal rails carry the mechanical load, while a separate insulated conductor system supplies grid power to the transporter or construction machine through a downward-facing slot. Segmenting the conductor and energizing only the section currently occupied by a machine is also exactly the direction I would pursue.

The architecture I would investigate

I would actually put **several conductors inside one protected enclosure**, rather than one red conductor:

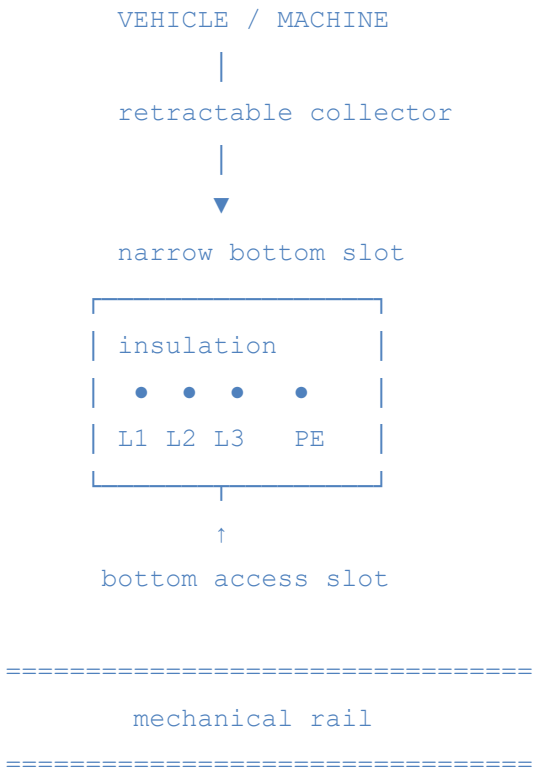

For a three-phase AC system you normally need **L1, L2 and L3**, plus a protective-earth arrangement. I would not intentionally use the structural running rails as the normal power return path. The entire conductor-rail assembly could sit perhaps a meter or so beside the transportation rails and low enough that it doesn't interfere with the 12 m × 24 m railcars.

The downward-opening groove in your drawing is a good feature because it makes accidental contact much more difficult and helps keep falling objects away from the conductor.

Why I like the idea

The biggest benefit is that your automated construction machinery no longer needs enormous diesel engines, batteries, or trailing multi-megawatt cables.

A machine doing excavation, concrete production, pumping, heavy lifting, welding, tunnel construction, or transporting a 12 × 96 m platform could simply pull power from the site grid.

That gives you: `grid → conductor rail → machine`

instead of: `grid → battery charger → huge battery → machine`

or: `diesel delivery → engine → machine.`

For machines operating repetitively on a fixed industrial site, the conductor rail could be particularly attractive.

It also fits the philosophy of the whole project: you are spending money once on a standardized infrastructure system and then using it for **hundreds of different machines and construction operations over decades**.

Your 0.1–10 MW range is reasonable, but I would split it into voltage classes

The upper end is where medium voltage becomes especially attractive.

For three-phase power:

So, approximately:

Machine load	Voltage	Approx. current*
0.1 MW	480 V	~127 A
0.5 MW	480 V	~633 A
1 MW	480 V	~1,270 A
1 MW	690 V	~880 A
2 MW	3.3 kV	~370 A
5 MW	6.6 kV	~460 A
10 MW	6.6 kV	~920 A
10 MW	7.04 kV	~865 A

*Assuming roughly 0.95 power factor.

That illustrates why I wouldn't try to supply a 5–10 MW excavator, crane, concrete system, or heavy transport machine at 440–690 V. The conductors and current collectors become enormous.

For the largest construction machines, **3.3–7.2 kV class distribution makes much more sense**. Medium-voltage industrial contactors in the roughly 7–12 kV class are established technology; for example, commercial vacuum contactors are available up to 12 kV.

I might therefore standardize the yard around something like:

690 V three-phase: small/medium machines.

6.6 kV three-phase: large construction/transport machinery.

The exact Chinese standard voltages would need to be selected during engineering.

Segmentation is one of the strongest parts of your idea

I agree that **most of the rail should normally be dead**.

Imagine the track divided into, say, 25–100 m electrical zones:

```
Segment 1 | Segment 2 | Segment 3 | Segment 4 | Segment 5
   OFF         OFF          ON          OFF         OFF
                             ▲
                          MACHINE
```

A machine entering Segment 3 communicates its identity and position to the site's power-control system.

Then:

- System confirms the machine is authorized.
- Collector is correctly engaged.
- Adjacent segment doors/access points are secure.
- Ground/insulation tests pass.
- Segment 3 is energized.
- Segment 2 is de-energized behind it as appropriate.
- Segment 4 can be energized just before transition.

This resembles the logic of a **moving electrical envelope** around the machine.

That means thousands of meters of conductor rail could exist while perhaps only **50–200 meters are energized around active equipment**.

I would not use ordinary mechanical relays

Your conceptual idea of switching sections with relays is correct, but at these powers I would use **industrial switchgear**.

At 6.6–7.2 kV and hundreds of amps, I'd expect equipment such as:

**vacuum contactors / vacuum circuit breakers

1. protection relays
2. ground-fault detection
3. isolation disconnects
4. visible grounding switches **

Commercial medium-voltage vacuum contactors are specifically designed for repeated industrial switching and are available up to 12 kV. IEC railway standards similarly treat high-power traction switching as dedicated circuit-breaker/disconnector/switchgear functions rather than ordinary low-voltage relays.

So, your "relay" becomes a **sectional power cabinet** every so many meters.

The warning lights are a good idea—but should be only one layer

I like your suggestion enormously from a human-factors standpoint.

When a segment becomes energized:

> bright red/orange LED strip illuminates continuously along that particular segment.

When dead:

> green or no illumination.

I would probably make the rail itself visibly communicate status:

> RED = ENERGIZED — KEEP CLEAR

But the light can't be the safety system.

Arc-flash and shock hazards can be serious even at ordinary low voltages, and OSHA emphasizes that even 120/208-V equipment can produce dangerous arc flashes. At several kilovolts and multi-megawatt available fault power, the electrical design must presume that an energized conductor is potentially lethal and use physical guarding, isolation and interlocking.

I'd therefore use at least:

**physical enclosure

- downward inaccessible opening
- segment isolation

- position detection
- collector interlocks
- ground-fault protection
- short-circuit protection
- emergency stop
- visible warning lights
- lockout/tagout capability.**

I would make the blue "stem" retractable

Your blue stem is another good concept, but I'd turn it into a sophisticated **collector assembly**.

Something like:

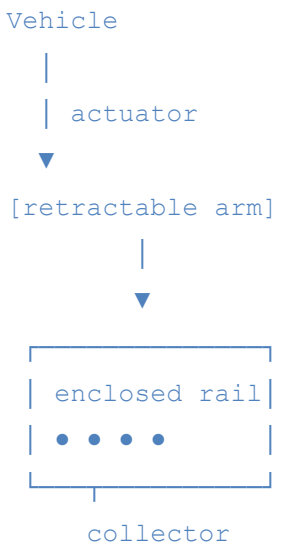

When a machine isn't using rail power, the collector retracts completely.

To connect:

> vehicle stops → alignment verified → collector extends → contact verified → electrical segment energizes.

For movement along the track, spring-loaded sliding shoes could maintain contact.

There should probably be **two collector sets separated longitudinally**, so a machine can transition smoothly between electrical segments.

For example:

```
MACHINE
```

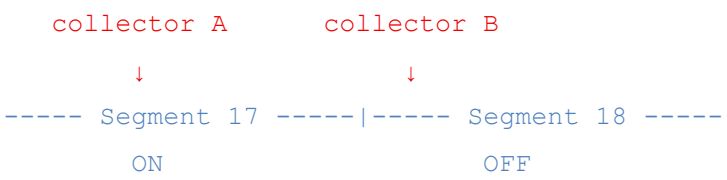

As B reaches Segment 18:

1. system verifies B contact;
2. Segment 18 energizes;
3. machine crosses boundary;
4. A clears Segment 17;
5. Segment 17 de-energizes.

That avoids interruption.

Water and dirt are the biggest problems with the ground-level concept

The environment at your nuclear/chemical construction site would be terrible for electrical contacts:

rain
salt spray near the waterfront
mud
concrete slurry
steel fragments
dust
snow/ice in some climates
chemicals
construction debris.

Therefore, I would make your green enclosure substantially more sophisticated than the simple U shape in the sketch.

Something closer to:

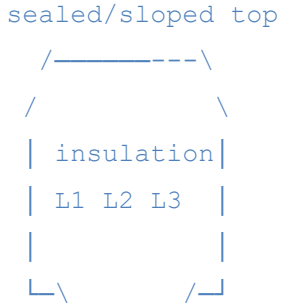

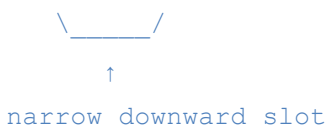

The slot could contain flexible insulating lips or overlapping barriers.

The conductor surfaces themselves could be **up inside the housing**, not immediately adjacent to the opening.

The blue collector stem reaches upward/sideways once inside the slot.

This gives you a "labyrinth" path, making it hard for fingers, tools, rain or debris to reach the conductors.

I would seriously consider placing it between the two mechanical rail systems

Suppose the enormous transporter runs on four rails:

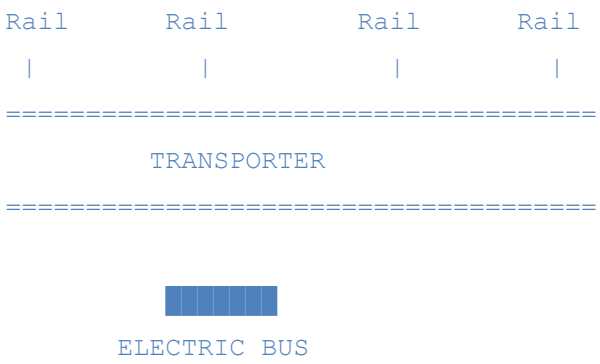

The conductor system could occupy a protected central region under the vehicle rather than an outer edge where workers frequently walk.

That has an important safety advantage:

> when a 1,000-ton machine is sitting over the energized conductor, humans physically cannot get to it easily.

When the machine leaves, the segment shuts off.

That might be better than putting an energized rail beside a pedestrian-accessible track.

A still safer alternative: recessed trench conductor

You might go one step beyond your picture:

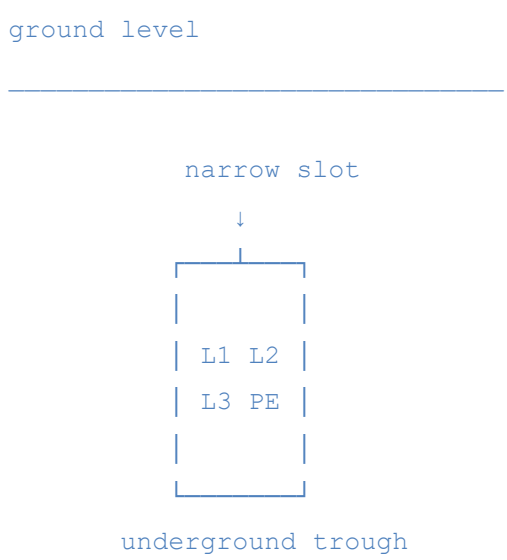

Now virtually the entire electrical system is **below grade**.

The collector descends through the slot.

Advantages:

- harder for workers to touch;
- protected from vehicle collision;
- doesn't consume surface space;
- can be drained;
- easier to place heavy protective concrete around it.

Disadvantages:

- drainage becomes critical;
- more excavation;
- harder maintenance;
- debris can enter slot.

I would prototype **both the above-ground enclosed rail from your drawing and a recessed version**.

Another useful feature: machines could regenerate energy

This gets particularly interesting with enormous moving masses.

Suppose your 12 × 96 m platform transporter is braking.

Its traction motors can operate as generators:

moving platform → motors → conductor rail → site electrical bus.

Similarly:

crane lowering 1,000-ton equipment → regenerative drive → grid

or:

excavator lowering boom → regeneration.

That isn't a revolutionary amount of energy relative to the nuclear plant, but it reduces braking heat and improves machine efficiency.

It could eliminate very large onboard power systems

Think about an automated tank-building machine.

Instead of putting:

diesel engine + generator + fuel tank

on the railcar, you put:

6.6-kV collector
→ transformer / converter
→ DC bus
→ motors

The equipment could then power:

- 5-MW excavation systems;
- 2-MW concrete plants;
- giant pumps;
- crushers;
- tunnel-boring equipment;
- automated welding;
- large gantry cranes;
- platform transporters.

Because the nuclear construction site itself has **gigawatts of electrical infrastructure**, supplying 5 or 10 MW to a construction machine is relatively modest from the site's perspective.

This might become a common interface for every automated construction machine

This is the part I like most.

You already have a proposed standard:

> 12 m × 24 m double-width railcar.

Add an electrical standard:

> Automated Construction Power Interface

For instance:

Mechanical:	four-rail / double-rail transporter geometry
Large-power interface:	6.6-kV 3-phase conductor rail
Auxiliary power:	690/400 V locally transformed
Control:	fiber + wireless network
Safety:	redundant segment-position interlock
Data:	machine ID, permitted power, rail segment status
Grounding:	dedicated PE/bonding conductor
Maximum unit power:	perhaps 10 MW.

Then every automated construction machine designed for the program has the same electrical pickup.

That would mean the same rail system that moves a chemical-processing platform can power:

> the transporter carrying it,
> the excavator preparing its bay,
> the concrete machine building its foundation,
> the crane placing it,
> the welding machines connecting it,
> and later the transporter removing it.

My main design change to your sketch

So, I would retain about **80% of your concept**, but change the electrical detail from:

> one insulated red rail + one blue collector

to:

> an enclosed, segmented 3-phase conductor-bus system with a dedicated protective conductor, a retractable multi-contact collector, and medium-voltage sectional switchgear.

And for multi-megawatt construction machinery, I would make something around **6.6 kV** the primary candidate rather than 440 V.

Your proposed segmentation and illuminated energized-zone indication are both particularly good ideas. The segmentation means this doesn't need to behave like an ordinary railway third rail where kilometers may remain continuously live; it can behave more like **an automatically moving power socket that exists only underneath the machine that currently needs it**.

Different Ways to Make Ammonia

Prompt to AI: What are the different methods for making liquid ammonia with heat and electricity from automated construction HTR-PM, and how much does each of these cost in units of $ per tonne of ammonia? Apply the same question to automated construction PWR. Assume: (a) we are looking for lowest cost approach with waste heat recovery, (b) heat pumps reduce cost, (c) we are in China, and (d) we are working with factory made platforms of equipment. Also, what is the current cost of making Ammonia in China using fossil fuel? Provide a summary table showing the different fossil fuel-based methods and the different green (less CO_2 emissions) methods. Include: (a) cost per tonne of ammonia, (b) tonnes of CO_2 emitted into atmosphere per tonne of ammonia.

Reference AI Conversation

Discussion, Electricity and Ammonia Cost (google drive, AI Conversation)

Response From AI:

Using the assumptions in your documents, and updating the earlier ammonia calculations to the new the current companion site-construction proposal electricity/heat costs, I get a somewhat different result from the older discussion:

The lowest-cost nuclear route is probably high-temperature steam electrolysis (HTSE/SOEC) + Haber–Bosch for both reactor types. For a mature Chinese factory-platform system, I would now use roughly $360–380/tonne of liquid NH_3 for the automated PWR and $400–420/tonne for the automated HTR-PM as central planning numbers. With particularly successful waste-heat recovery, industrial heat pumps, high utilization and repeated platform manufacture, the lower ends could be around $350/t for PWR and $390/t for HTR-PM.

That puts the automated PWR surprisingly close to present Chinese coal-ammonia economics.

Why the new estimates are lower

The current companion site-construction proposal nuclear document now uses approximately $37/MWh-electric + $15/MWh-thermal for automated HTR-PM, versus $31/MWh-electric + $10/MWh-thermal for automated PWR.

The ammonia discussion gives an LWR-integrated HTSE requirement of approximately 6.84 MWh-electric plus 1.13 MWh-thermal per tonne of NH_3, including the major downstream electricity requirement. That means the nuclear-energy contribution alone becomes:

HTR-PM: $6.84 \times \$37 + 1.13 \times \$15 \approx$ **$270/t NH_3** (heat and electricity cost only)

PWR: 6.84 × \$31 + 1.13 × \$10 ≈ **\$223/t NH_3** (heat and electricity cost only)

The remainder is electrolyzer capital and stack replacement, air separation, compression, Haber–Bosch, refrigeration, water treatment, storage, maintenance and financing. Your earlier HTR-PM model put those non-energy items at roughly \$120–190/t; the detailed PWR model put them at roughly \$130–175/t.

The factory-platform concept is particularly relevant here because it explicitly proposes repetitive factory manufacture of hydrogen and ammonia modules, shared utilities, air separation, refrigeration/storage and standardized nuclear interfaces rather than designing each chemical plant individually. The ammonia work package also explicitly calls for **Haber–Bosch heat recovery and recycle**.

Nuclear ammonia pathways

Here are the pathways I think are worth carrying in the economic model.

Category / reactor	Hydrogen pathway → NH_3	Estimated liquid NH_3 cost, China ($/t)	Direct CO_2 emitted (t CO_2/t NH_3)	Assessment
Fossil — China coal	Coal gasification → H_2 → Haber–Bosch	$325–370; **central ≈ $348**	~4.0–4.8 "fossil fuel"	Current Chinese incumbent; cheapest benchmark
Fossil — natural gas	SMR → H_2 → Haber–Bosch	~$350–500*	~1.8–2.6 "fossil fuel"	Lower emissions than coal, but gas is generally less attractive in China
Renewable electricity	Wind/solar → alkaline / PEM → Haber–Bosch	~$400–820; RMI ≈ $705	≈0 direct "green"	Commercial technology; cost dominated by electricity and storage
Automated HTR-PM	HTSE/SOEC → Haber–Bosch, heat recovery	$390–430; **central ≈ $405**	≈0 direct "green"	Best HTR-PM route
Automated HTR-PM	Alkaline electrolysis → Haber–Bosch	~$550–600	≈0 direct "green"	Simpler/mature electrolyzer, but much more electricity
Automated HTR-PM	PEM electrolysis → Haber–Bosch	~$580–630	≈0 direct "green"	Flexible, but generally higher electrolyzer cost
Automated HTR-PM	Sulfur–iodine, 38%-class cycle + electric topping → Haber–Bosch	~$500–620	≈0 direct "green"	Technically possible; too complex to beat HTSE in present HTR-PM
Automated HTR-PM	Advanced ~50% sulfur–iodine + strong recuperation/topping	~$450–550	≈0 direct "green"	Interesting R&D target, not commercially demonstrated
Automated PWR (nuclear)	HTSE/SOEC → Haber–Bosch, optimized heat recovery	$350–390; **central ≈ $370**	≈0 direct "green"	Lowest-cost nuclear route
Automated PWR	Alkaline electrolysis → Haber–Bosch	~$490–530	≈0 direct "green"	Technically straightforward but electricity-intensive
Automated PWR	PEM electrolysis → Haber–Bosch	~$520–570	≈0 direct "green"	Similar issue; little reason to prefer it for baseload NH_3
Automated PWR	Sulfur–iodine + massive electric topping	Probably >$600	≈0 direct "green"	Not economically sensible; PWR temperature is far too low

*The natural-gas number is a rough China planning range rather than a current Chinese plant quotation. Published international natural-gas ammonia production estimates are commonly around $250–450/t; China's relatively expensive gas generally pushes the economics toward the high end.

The CO_2 column is direct/process CO_2 emitted at the ammonia plant, not full lifecycle CO_2-equivalent. Nuclear and renewable routes have essentially no direct carbon-containing feedstock or fuel combustion in this configuration, so direct CO_2 is near zero; their lifecycle emissions are small but

nonzero and should be calculated separately. Your industrial-platform proposal appropriately says that nuclear ammonia should not simply be labeled “green” without publishing lifecycle boundaries. A Chinese study likewise treats fully renewable electrolysis ammonia as zero direct production-process CO_2 when the full energy chain is green.

For fossil ammonia, recent assessments emphasize just how important the feedstock is. China remains overwhelmingly coal-oriented in ammonia production; a Columbia analysis cites about 3.5–4.5 t CO_2/t NH_3 for coal ammonia, versus roughly 1.8–2.6 t/t for natural-gas ammonia. Your own discussion used about 4.6 t CO_2/t NH_3 as the representative Chinese coal number, so I would use 4.6 t/t centrally in your business-plan calculations.

Why HTSE is the winner for both reactors

Your industrial-processing document already identifies the right competition: HTSE versus alkaline and PEM electrolysis. Low-temperature alkaline/PEM requires roughly **10 MWh-electricity per tonne of NH_3**, while nuclear HTSE needs roughly 6.8–7.0 MWh-electric plus about 1.1 MWh-thermal.

At $37/MWh HTR-PM electricity, 10 MWh alone costs roughly **$370/t** before buying the electrolyzers, air separation, Haber–Bosch plant, refrigeration and storage. For the PWR it is about **$310/t**. That is why alkaline/PEM do not beat HTSE even though their technology maturity is attractive.

The more interesting result is that PWR HTSE beats HTR-PM HTSE even though HTR-PM supplies much hotter heat. Most of the monetary input to HTSE ammonia is still electricity. The INL configuration discussed in your file was actually developed for a light-water reactor: PWR steam around 250°C makes and preheats steam and electrical topping raises it to SOEC temperature. Consequently, saving about $6/MWh on electricity with the PWR is worth substantially more than the HTR-PM's superior heat quality for this particular product.

That does not make the HTR-PM inferior generally. It remains much more attractive for chemical processes that actually consume large quantities of 300–500+°C process heat. The distinction in the current companion site-construction proposal—approximately 520°C-class HTR-PM steam versus sub-300°C PWR steam—is important for those industries.

What waste-heat recovery and heat pumps should do

I would design both ammonia plants around aggressive heat integration, but I would **not** assume that a heat pump magically supplies the 800°C SOEC or 900°C sulfur–iodine requirement.

The valuable opportunities are lower-temperature ones: recover Haber–Bosch reaction heat to make steam; recover compressor/intercooler and ammonia-condenser heat; use mechanical-vapor recompression or industrial heat pumps to turn perhaps 60–100°C waste heat into 120–200°C process

heat; preheat electrolysis water and steam; drive evaporators and water treatment; and reduce refrigeration and auxiliary-steam loads. Your discussion specifically identifies these uses and cautions against trying to raise low-grade waste heat all the way to electrolysis temperature.

Some of that heat recovery is already embedded in the underlying HTSE and ammonia cost assumptions, so I would not subtract another huge amount. I would treat optimized waste-heat recovery, heat pumps, common utilities and repeated factory construction as what moves the PWR from roughly **$375–400/t toward $350–370/t**, and the HTR-PM from roughly **$420/t toward $390–410/t**.

Why sulfur–iodine probably does not win

For the existing HTR-PM, sulfur–iodine has a temperature mismatch. HTR-PM reactor outlet temperature is around 750°C and useful steam is roughly 500°C, whereas sulfuric-acid decomposition in the S–I cycle needs roughly **850–900°C**. It therefore requires electrical topping or a future hotter reactor.

The 38.1%-efficient S–I reference case also consumes roughly 18.2 MWh-thermal per tonne of ammonia just to produce the necessary hydrogen. At $15/MWh HTR-PM heat that is about **$274/t for reactor heat alone**, before the very substantial acid/iodine plant, high-temperature heat exchangers, nitrogen separation, Haber–Bosch, refrigeration and storage. Strong internal recuperation helps, but the plant is much more complicated than an SOEC platform. Hence my **~$500–620/t** present-HTR-PM estimate.

A 50%-efficient future S–I cycle could lower that materially, potentially toward **$450–550/t**, but that is an R&D outcome rather than something I would put into the central economic case.

For a PWR, sulfur–iodine makes even less sense because the reactor cannot directly provide the high-temperature heat. The PWR should simply use HTSE.

What does fossil ammonia actually cost in China today?

There is an important distinction between **production cost** and **market price**.

RMI's detailed China ammonia analysis uses **RMB 2,200–2,500/t** as the conventional coal-ammonia cost range and **RMB 2,350/t** as its central benchmark. At the August 8, 2026 exchange rate of roughly RMB 6.75/$, that corresponds to approximately **$326–371/t**, with a central value of about **$348/t**.

The actual market is currently slightly lower. On August 7, 2026, the reported Chinese liquid-ammonia domestic spot price was about **RMB 2,203/t**, or roughly **$327/t**. A separate June 2026 China FOB benchmark was **$339/t**.

So, for your documents, I would use:

> Current Chinese coal ammonia production cost: ≈ $350/t, with a reasonable $325–370/t range.
>
> Current Chinese ammonia market price: ≈ $325–340/t.

The main result

The comparison I would carry forward in the Climate Moonshot work is therefore:

> Coal incumbent: ~$350/t and ~4.6 t CO_2/t NH_3.
>
> Automated PWR + factory-platform HTSE: ~$365/t and ≈0 direct CO_2/t.
>
> Automated HTR-PM + factory-platform HTSE: ~$405/t and ≈0 direct CO_2/t.

That is substantially more interesting than the older $450 HTR-PM / $400 PWR comparison in the Discussion document, because the current companion site-construction proposal has lowered the assumed automated-reactor LCOEs to $37 and $31/MWh respectively. The implication is that an optimized automated-PWR/HTSE/ammonia super-site could conceivably get within a few tens of dollars per tonne of today's Chinese coal ammonia without a carbon price. It would not necessarily beat the cash marginal cost of an already-built coal-ammonia plant, but it could become very competitive for new ammonia capacity while eliminating roughly 4–5 tonnes of direct CO_2 for every tonne of ammonia produced.

The biggest technical/economic number to validate in the proposed $1B platform fund is therefore not sulfur–iodine efficiency. It is whether repeated Chinese factory-made **HTSE + air-separation + Haber–Bosch + heat-recovery platforms can actually achieve roughly $130–160/t of non-energy cost at very high utilization**. If they can, the automated-PWR case becomes particularly compelling.

Ships Powered by Ammonia

Prompt to AI: What are the different ways to power large ships with liquid ammonia? Assume we use high temperature fuel cells on ships to convert ammonia to electricity, to power electric motors. What is the temperature of the fuel cell with each method? What are the advantages and disadvantages of each method? Assume: (a) we do what we can to lower cost, (b) we do waste heat recovery, (c) we use heat pumps to lower cost if we can, and (d) we filter/scrub the exhaust as needed to keep pollutants at an acceptable level.

Response From AI:

Yes. If the design constraint is **liquid ammonia → high-temperature fuel cell → electricity → electric propulsion motors**, there are several distinct architectures. For a very large ship, I would concentrate on **SOFCs** today and keep **protonic ceramic fuel cells (PCFCs)** as a potentially lower-cost next-generation option. Conventional oxide-ion SOFCs generally operate around 500–1,000°C, while proton-conducting ceramic cells operate roughly 400–700°C.

Main ammonia-to-electricity options for a large ship

#	Architecture	Fuel-cell temperature	What happens to NH_3	Relative attractiveness
1	Direct-ammonia SOFC, internal cracking	~700–850°C	NH_3 enters SOFC and cracks to H_2 + N_2 inside/near anode	Potential lowest CAPEX / simplest
2	Partial external cracker + SOFC	~650–800°C	Some NH_3 cracked upstream; rest internally cracked	My preferred near-term design
3	Full external cracker + SOFC	~650–800°C	Nearly all NH_3 → H_2 + N_2 before SOFC	Very robust, somewhat higher CAPEX
4	External cracker + PCFC	~500–650°C	NH_3 cracked first; H_2/N_2 feeds protonic cell	Very promising future low-cost design
5	Direct/integrated ammonia PCFC	~550–650°C	Cracking catalyst integrated tightly with fuel-cell module	Potentially excellent, but immature
6	Cracked NH_3 + molten-carbonate fuel cell	~600–700°C	NH_3 cracked to H_2; MCFC converts H_2	Technically possible, not my choice

DOE gives representative temperature ranges of 500–1,000°C for oxide-ion solid-oxide cells, 400–700°C for proton-conducting ceramic cells, and 600–700°C for molten-carbonate cells. For ammonia specifically, experiments and models increasingly focus on SOFC operation around 650–800°C because ammonia decomposition becomes much easier as temperature rises; studies report significant degradation/performance penalties below about 700°C in some direct-ammonia SOFC designs.

1. Direct-ammonia SOFC: probably the simplest system

The liquid NH_3 is pumped, vaporized and preheated. Instead of building a separate hydrogen-production plant aboard ship, gaseous NH_3 goes essentially straight into the hot SOFC fuel side:

Liquid NH_3 → vaporizer → recuperator → SOFC → DC electricity → inverter → motor

At roughly **700–850°C**, ammonia readily decomposes:

$2\ NH_3 \rightarrow N_2 + 3\ H_2$

The hydrogen is then electrochemically consumed in the SOFC. Research on direct-ammonia SOFCs commonly operates in approximately the 600–800°C region, and performance becomes much more attractive around 700°C and above.

Advantages: It eliminates a large standalone ammonia cracker and potentially eliminates hydrogen purification, storage and compression. The cracking reaction is endothermic, so internal ammonia decomposition actually absorbs heat inside the hot fuel-cell system, which can help thermal management. There are no carbon atoms in the fuel, so the conversion itself produces no CO_2.

Disadvantages: The fuel-cell stack becomes both an electrochemical device and part of the ammonia cracker. That complicates temperature control. Too much ammonia cracking at one location can create cold spots and thermal gradients. Direct NH_3 exposure can also affect conventional nickel-containing anodes and accelerate degradation. Research specifically continues on catalysts and anode materials to improve direct-ammonia SOFC durability.

This could eventually be the **lowest-equipment-cost solution**, but I would be reluctant to make the first 50–100 MW ship propulsion system depend entirely upon internal cracking.

2. Partial external cracking + SOFC: my preferred architecture

I think this is the most interesting configuration for the kind of cost-minimized system you are envisioning.

Instead of choosing between "all cracking outside" and "all cracking inside," split the job:

Liquid NH_3 → vaporizer → waste-heat-heated cracker → partially cracked $NH_3/H_2/N_2$ → SOFC → electric motor

The cracker could operate around **550–700°C**, while the SOFC might operate around **700–800°C**.

You deliberately **do not need to manufacture pure hydrogen**. The SOFC can consume hydrogen mixed with nitrogen, and residual ammonia can finish decomposing in the hot fuel-cell module.

That eliminates hydrogen separation and purification equipment.

The big advantage is thermal management. You can adjust the external cracker so that, say, enough ammonia has been cracked to prevent excessive cooling or ammonia exposure inside the SOFC, while allowing some internal cracking to absorb surplus stack heat.

A modeling study comparing direct feed with external cracking found the performances become quite similar at sufficiently high SOFC inlet temperatures, illustrating why there may be little reason to pay for complete external conversion if the stack itself is already hot enough to finish the job.

Advantages: Less expensive than a complete hydrogen plant; better SOFC temperature uniformity than pure internal cracking; lower direct ammonia exposure of the anode; no hydrogen purification; extremely good integration with waste heat; relatively easy to divide into standardized modules.

Disadvantages: You add cracker catalyst beds, heat exchangers, valves and pressure drop. There is still some interaction between ammonia cracking and stack thermal behavior.

For a large standardized ship, **this is the architecture I would develop first.**

3. Completely crack ammonia before the SOFC

The conservative alternative is:

NH_3 → external catalytic cracker → $H_2 + N_2$ → SOFC

The SOFC remains around **650–800°C**, while the ammonia cracker could operate around **600–700°C** with an inexpensive nickel-based catalyst system. Higher-performance catalysts can permit lower cracking temperatures, but cost and catalyst lifetime become important.

There is usually **no reason to separate the nitrogen from the hydrogen**. Feed the cracked H_2/N_2 mixture directly to the SOFC.

That is important economically: you are building an ammonia cracker, not a complete ammonia-to-pure-hydrogen plant.

Advantages: The stack sees little residual NH_3, making stack materials and controls easier. Cracker conditions can be optimized independently. SOFC thermal gradients are reduced. Fuel quality becomes predictable. Replacement of cracker catalyst is much easier than replacement of damaged fuel-cell stacks.

Disadvantages: More hardware, more heat exchangers, more catalyst, greater volume and mass and somewhat higher capital cost.

For an early commercial ammonia-fuel-cell ship, I would expect this approach to be easier to certify and operate than pure direct-ammonia operation.

4. External ammonia cracker + PCFC

This is particularly interesting longer term.

A **protonic ceramic fuel cell** performs a similar basic job to an SOFC but transports protons through its ceramic electrolyte rather than oxygen ions. That allows substantially lower operating temperature—typically approximately **400–650°C** rather than conventional SOFC temperatures.

A practical marine arrangement could be:

NH_3 → ~600°C cracker → H_2/N_2 → ~550–600°C PCFC → electricity

Experiments have demonstrated ammonia-fed PCFC operation at **600°C for more than 1,200 hours** when an upstream ammonia-cracking catalyst was used, and maximum power density reached 877 mW/cm² at 650°C in that experiment.

There is another thermodynamic advantage. In an oxide-ion SOFC, water is formed on the fuel side. In a PCFC, hydrogen is converted to protons and water is primarily produced on the air side. That helps avoid diluting the remaining hydrogen fuel with its own water product.

Advantages: Lower-temperature alloys and seals can potentially be used; less thermal stress; faster startup; potentially longer component life; potentially cheaper balance of plant; good efficiency; and a compact cracker/fuel-cell combination around a common ~600°C thermal level.

Disadvantages: PCFC technology is nowhere near the commercial maturity of conventional large SOFC technology. Large marine-scale stacks have not yet accumulated the operating experience that one would want for a propulsion plant. DOE categorizes several such next-generation fuel-cell approaches as earlier-stage R&D.

For a development program with a **10–20 year horizon**, this is the technology I would watch most closely.

5. Direct/integrated ammonia PCFC

You could eliminate the separate cracker again:

NH_3 → catalyst integrated into PCFC module → electricity

Target temperature would probably be around **550–650°C**.

This sounds ideal—lower temperature than an SOFC and no large cracker—but the materials problem is harder. In one detailed study, uncracked ammonia caused substantial degradation in the protonic ceramic material, whereas putting an effective ammonia-cracking catalyst immediately upstream greatly improved durability.

So there is a semantic distinction between a "direct ammonia PCFC" and what I would actually build. I would place the ammonia-decomposition catalyst **inside the same factory-built module but immediately upstream of the electrochemical cells**.

That gives almost the packaging advantage of direct ammonia while protecting the ceramic stack.

6. Molten-carbonate fuel cell

MCFCs operate around **600–700°C**. Ammonia could be cracked and the resulting hydrogen supplied to an MCFC.

I would not pursue this for an ammonia ship.

The fundamental nuisance is that a molten-carbonate cell needs **CO_2 at the cathode** to transport carbonate ions. On a natural-gas system this can be integrated into the carbon-containing exhaust stream. With carbon-free ammonia, you now need to supply and recycle CO_2 merely to make the fuel cell work.

It throws away one of ammonia's biggest advantages: a carbon-free fuel loop.

Waste heat changes the economics substantially

For these systems, I would **not** regard the 600–800°C exhaust as "waste" in the ordinary sense.

The thermal train should probably look like:

SOFC exhaust → ammonia cracker → incoming-air heater → incoming-NH_3 vaporizer/preheater → steam/ORC or other bottoming cycle → ship heating/hot water → final heat rejection

High-temperature fuel cells are particularly attractive precisely because their hot exhaust permits this kind of cogeneration. DOE notes that fuel-cell combined heat-and-power systems can achieve very high overall energy utilization.

For a giant ship, I would investigate going one step further:

SOFC + gas turbine generator

The hot pressurized exhaust drives a turbine connected to another generator. Published marine SOFC/gas-turbine studies report system efficiencies around or above 60% in modeled designs, and more aggressive ammonia-SOFC system analyses have projected electrical efficiencies approaching 70%; these are model results, not demonstrated commercial ship performance.

That electricity simply joins the ship's DC bus and drives the same electric propulsion motors.

Heat pumps are much less valuable here than in the ammonia factory

This is an important difference from the ammonia-production system we discussed previously.

A heat pump is useful when you have, for example, **70°C waste heat and need 150°C process heat**.

On the ship, you already have enormous quantities of waste heat at **500–800°C**.

It would make little economic sense to use electricity to run a heat pump to make the ~600°C heat required by an ammonia cracker. Use a recuperative heat exchanger directly.

Heat pumps can nevertheless help with the **low-temperature end of the ship**: HVAC, domestic hot water, refrigerated ammonia-tank management, cooling systems and recovery of low-grade heat after the high-temperature duties have been satisfied. But they are a secondary optimization rather than a key element of the propulsion architecture.

What comes out of the exhaust?

Ideally, ammonia through a fuel cell ultimately becomes essentially:

$NH_3 + O_2 \rightarrow N_2 + H_2O$

so there is **no combustion-derived CO_2**.

However, a real ship needs to control three nitrogen pollutants:

NH_3 slip, NO_x, and N_2O.

Fuel cells themselves avoid the very high flame temperatures responsible for much combustion NO_x, and an experimental direct-ammonia PCFC has demonstrated ammonia operation at 600°C without measured NO_x under those test conditions. Nevertheless, residual ammonia and hydrogen in the

anode exhaust normally need to be treated or recycled; if they are burned in an afterburner, nitrogen-containing species can be produced.

I would therefore design the propulsion package from the beginning with **fuel recycle + ammonia-slip catalyst + NO_x treatment + N_2O catalyst as necessary**, rather than treating emissions control as an afterthought.

IMO already has interim safety guidelines covering both ammonia-fueled ships and shipboard fuel-cell power installations.

What I would build

For a first cost-optimized large ocean-going ship, my reference design would be:

```
Liquid NH₃ tanks
→ pump
→ vaporizer using waste heat
→ ~600–700°C inexpensive catalytic partial cracker
→ ~700–800°C SOFC modules
→ DC bus
→ variable-frequency inverters
→ electric propulsion motors
```

and thermally:

```
hot SOFC exhaust
→ ammonia cracker
→ fuel/air recuperators
→ optional bottoming-cycle generator
→ low-temperature ship heat/refrigeration recovery
→ exhaust treatment
```

I would probably **partially rather than completely crack the ammonia**. This minimizes cracker size and cost while preventing the SOFC stack from having to perform all of the endothermic ammonia decomposition itself.

For the longer-term system, I would aggressively develop:

```
~600°C cracker + ~550–600°C PCFC.
```

If PCFC durability, manufacturing and MW-scale stack costs mature, that could eventually become **better than the 750°C-class SOFC** because the entire machinery plant becomes cooler, cheaper and easier to insulate while preserving the fundamental efficiency advantages of a ceramic fuel cell.

So, if we are designing your factory-made industrial/ship equipment concept, I would carry three serious competing designs through R&D: direct/internal-cracking SOFC, partial-cracking SOFC, and externally cracked PCFC. My expectation is that partial-cracking SOFC wins first deployment, while PCFC has the best chance of winning long-term.

Ammonia to Electricity Conversion Efficiency (e.g. on a ship)

Prompt to AI: What efficiencies do you expect with your "three serious competing designs" assuming you do what you can to maximize efficiency? Assume we can get liquid green ammonia at $365/tonne via Automated China PWR (your estimate). How much does rotational energy on a ship cost using your 3 fuel cell methods and $365/tonne green ammonia ($ per MWh-mechanical energy)? How does this compare to shipping high sulfur fuel oil? low sulfur fuel oil?

Response From AI:

Yes. At $365/tonne for liquid green ammonia, the economics become quite interesting. On a fuel-only basis, an optimized ammonia/SOFC propulsion system can come close to today's high-sulfur fuel oil and can be materially cheaper than today's very-low-sulfur fuel oil.

I would define efficiency here as ammonia LHV energy in the tank → net rotating mechanical energy at the propulsion-motor shaft, after ammonia pumps/cracking, blowers, power electronics, exhaust treatment and motor losses. Waste heat only counts toward efficiency when we actually convert it back into useful propulsion electricity. The current IMO lifecycle calculation data use about 18.6 MJ/kg LHV for ammonia, or 5.17 MWh-thermal per tonne.

Fossil fuel used with ships:

HSFO: High sulfur fuel oil (definition)
VLSFO: Very low sulfur fuel oil (definition)

Fuel cells that convert ammonia to electricity:

PCFC: Protonic fuel cell (wiki)
SOFC: Solid oxide fuel cell (wiki)

My engineering estimates

DOE puts conventional SOFC electrical efficiency around 60%, while optimized direct-ammonia SOFC studies have also reported around 60% and some highly integrated modeled systems approach the upper-60% range. These higher numbers are engineering-model results, not demonstrated 50–100 MW ship systems. PCFC systems are less mature; one detailed PCFC system study predicted about 55.1% electrical efficiency, though the technology has advantages from intermediate-temperature operation.

For an aggressively optimized ship, I would therefore carry these planning numbers:

Propulsion architecture	Fuel-cell temperature	Expected net electricity efficiency	Net ammonia → shaft efficiency	Fuel cost at $365/t NH_3
Direct/internal-cracking SOFC	~700–800°C	~65–68% with strong WHR/bottoming	~64% central; ~62–67% range	~$110/MWh-mechanical
Partial external cracker + SOFC	Cracker ~600–700°C; SOFC ~700–800°C	~66–69%	~65% central; ~62–67% range	~$109/MWh-mechanical
External cracker + PCFC	Cracker ~550–650°C; PCFC ~500–600°C	~59–62% optimized	~58% central; ~55–60% range	~$122/MWh-mechanical
HSFO + large low-speed diesel + scrubber	Fossil Fuel Combustion	~50% shaft before scrubber load	~49% effective	~$107/MWh-mechanical at today's Singapore HSFO price
VLSFO + large low-speed diesel	Fossil Fuel Combustion	~50%	~50%	~$146/MWh-mechanical at today's Singapore VLSFO price

The one- or two-percentage-point distinction between the two SOFC cases should **not** be treated as established fact. They are essentially tied thermodynamically at this stage. I slightly favor partial cracking for an actual ship because it gives much better control of stack temperature, ammonia exposure and catalyst replacement, rather than because it has some large inherent efficiency advantage.

The ammonia calculation

At $365/t:

Fuel energy cost: $365 / 5.17 MWh_{th}/t = $70.6/$MWh_{th}$.

Therefore:

Direct SOFC, 64% shaft efficiency: $70.6 / 0.64 ≈ $110/MWh-shaft.

Partial-cracking SOFC, 65%: $70.6 / 0.65 ≈ $109/MWh-shaft.

PCFC, 58%: $70.6 / 0.58 ≈ $122/MWh-shaft.

Those results are almost entirely controlled by two numbers: the $365/t ammonia price and the fuel-to-shaft efficiency.

Comparison with today's bunker fuel

For a concrete benchmark, Singapore is useful because it is the world's major bunkering hub. Current Ship & Bunker prices are approximately **$588.50/t for IFO380/HSFO** and **$821.50/t for VLSFO**. These prices are unusually important to timestamp because bunker prices are volatile.

IMO uses approximately **40.4 MJ/kg LHV for heavy fuel oil**, equivalent to about **11.22 MWh/t**. Large low-speed two-stroke diesels can exceed 50% thermal efficiency.

Thus, HSFO at $588.50/t costs about $588.50 / (11.22 × 0.50) ≈ $105/MWh-shaft before scrubber energy.

> But HSFO requires an exhaust-gas cleaning system to comply with the global sulfur limit where such use is permitted; IMO recognizes scrubbers as a compliance route. With a roughly 2% scrubber energy penalty, the effective fuel cost is approximately $107/MWh-shaft.

VLSFO doesn't require the sulfur scrubber for the global 0.50% sulfur standard. At $821.50/t and 50% efficiency: $821.50 / (11.22 × 0.50) ≈ $146/MWh-shaft.

So, the striking result is:

Fuel/technology	Fuel cost per MWh of shaft energy		Relative to $365/t NH_3 partial-SOFC
HSFO + diesel + scrubber	~$107	(fossil fuel, high sulfur)	~2% cheaper
Green NH_3 + partial-cracking SOFC	~$109	(green ammonia)	Baseline
Green NH_3 + direct SOFC	~$110	(green ammonia)	~1% higher
Green NH_3 + PCFC	~$122	(green ammonia)	~12% higher
VLSFO + diesel	~$146	(fossil fuel, low sulfur)	~34% more expensive

At these particular August 2026 bunker prices, therefore, your **$365/t nuclear green ammonia is almost at fuel-cost parity with HSFO**, assuming we succeed in building a ~65%-efficient ammonia-to-shaft SOFC system.

And it is **substantially cheaper than VLSFO**: roughly $109 versus $146/MWh-shaft, a savings of about **25%**.

Waste-heat recovery is crucial

For the SOFC cases, I would design the thermal system very aggressively. SOFCs themselves are ordinarily around 60% electrical efficiency; DOE specifically highlights their high-temperature exhaust and suitability for hybrid gas-turbine cycles. The way I get from roughly 60% stack/system

performance into the mid-60s at the propulsion shaft is not a miraculous fuel cell—it is **system integration**.

The hot exhaust first provides the endothermic heat for ammonia cracking and NH_3 vaporization/preheating. Remaining high-grade heat can then support a gas-turbine or other bottoming generator. Lower-grade heat supplies ship services. Power electronics and a very large marine electric motor then impose only a few percentage points of loss.

Heat pumps are useful only near the low-temperature end. It would be counterproductive to run a heat pump to create 600°C cracker heat when there is already 700–800°C SOFC exhaust available. Heat pumps can instead recover low-grade heat for HVAC, hot water, refrigeration and other hotel loads, thereby leaving more generated electricity available for propulsion.

There is another important comparison

The numbers above compare **fuel cost**, not total propulsion cost.

That distinction could determine whether ammonia actually wins. An HSFO ship already has a relatively inexpensive, extraordinarily mature two-stroke diesel connected almost directly to the propeller shaft. The ammonia ship would need ammonia tanks and fuel handling, cracker equipment, large ceramic fuel-cell stacks, power electronics, electric motors, thermal-recovery equipment and emissions-control equipment.

On the other hand, the ammonia system eliminates a great deal of diesel-engine machinery and potentially has fewer major moving parts. The direct-SOFC route also avoids the large external cracker. Whether stack replacement overwhelms those savings is something that needs an actual $/kW and stack-life model.

There is also a **fuel-volume penalty**. Ammonia's volumetric energy density is much lower than HFO's. IMO analyses have estimated roughly 2.7 times the liquid-fuel volume per unit chemical energy for ammonia versus HFO. The ~65% fuel-cell efficiency versus ~50% diesel efficiency reduces that propulsion-level penalty substantially, but ammonia still takes considerably more tank volume.

What I think is most significant

The number I would put in your Climate Moonshot analysis is:

> Automated Chinese PWR → $365/t green NH_3 → optimized partial-cracking SOFC ship → approximately $109/MWh of mechanical propulsion energy.

Compare that with approximately:

HSFO today: ~$107/MWh-mechanical

VLSFO today: ~$146/MWh-mechanical

That means **$365/t green ammonia is already around the break-even fuel price against today's HSFO and comfortably below today's VLSFO**, provided the shipboard system really achieves approximately 65% ammonia-to-shaft efficiency.

The next question I would investigate is even more consequential: **including the capital cost and replacement cost of the fuel cells, crackers, electric motors and ammonia tanks, what is the total $/MWh-shaft for a 50-MW or 100-MW container ship over 25 years?** That would tell us whether $365/t nuclear ammonia merely has competitive *fuel* economics or could actually produce a cheaper ship to operate overall.

Cost of Ammonia Fuel Cell on Large Ships

Prompt to AI: Including the capital cost and replacement cost of the fuel cells, crackers, electric motors and ammonia tanks, what is the total $/MWh-shaft for a 50-MW or 100-MW container ship over 25 years? Compare this with traditional HSFO (High Sulfur Fuel Oil) and VLSFO (Very Low Sulfur Fuel Oil) approaches. Assume we are looking at new construction in all cases.

Response From AI:

Including the propulsion hardware, ammonia tanks, fuel-cell stack replacement, cracker, power electronics, electric motors, waste-heat recovery, emissions equipment, and O&M changes the comparison—but **$365/t ammonia still looks competitive with VLSFO**.

My central estimate is:

Partial-cracking SOFC + $365/t ammonia:	≈ $165/MWh-shaft	“green”
Direct-ammonia SOFC:	≈ $166/MWh-shaft	“green”
PCFC:	≈ $189/MWh-shaft	“green”
HSFO + diesel + scrubber:	≈ $153/MWh-shaft	fossil fuel
VLSFO + diesel:	≈ $183/MWh-shaft	fossil fuel

So, the optimized SOFC-ammonia ship comes out roughly **8% more expensive than HSFO, but 10% cheaper than VLSFO**, before putting any price on CO_2.

25-year levelized comparison

I modeled a new container ship for 25 years at a 7% real cost of capital. For utilization, I used the operating history in a published study of a 56.38-MW, 20,000-TEU containership. It had about 561 TJ/year of useful propulsion output, equivalent to about 156 GWh/year or a **31.6% full-power-equivalent utilization factor**.

For fossil fuel I retain the August 7 Singapore bunker prices from our previous comparison: **$588.50/t IFO380/HSFO and $821.50/t VLSFO**. I hold those prices constant in real dollars for the 25-year calculation; that is a comparison scenario, not a 25-year oil-price forecast.

New-build propulsion system	NH_3/fuel → shaft efficiency	Fuel	Initial capital recovery	Fuel-cell replacements	O&M	Total $/MWh -shaft
Direct/internal-cracking SOFC + NH_3	64%	$110	$38	$8	$10	$166
Partial-cracking SOFC + NH_3	65%	$109	$38	$8	$11	$165
External cracker + PCFC + NH_3	58%	$122	$45	$10	$12	$189
HSFO + low-speed diesel + scrubber	49% net	$108	$38	—*	$8	$153
VLSFO + low-speed diesel	50%	$147	$30	—*	$6	$183

*Diesel overhaul/rebuild costs are placed in O&M rather than modeled as periodic replacement of the entire engine.

The published container-ship study assumed 50% efficiency for HFO ICE and 60% for an ammonia-SOFC system; our 64–65% SOFC cases are the more aggressively optimized configurations we discussed, incorporating high-temperature heat recuperation and bottoming-cycle recovery. The study gives ammonia an LHV of 18.6 MJ/kg and HFO 40.2 MJ/kg.

Capital cost

This is where the comparison becomes interesting. Ammonia propulsion is not necessarily several times the capital cost of a conventional ship if high-temperature fuel cells reach mature mass-production costs.

A published marine SOFC study assumed **$900/kW for an ammonia-SOFC propulsion system**, versus $600/kW for an HFO diesel. It estimated ammonia-tank cost at $603/m³ and HFO tank cost at $883/m³. For its 56.38-MW vessel, the ammonia system required 32,760 m³ of tank capacity versus 15,142 m³ for HFO. DOE's SOFC development targets are also useful anchors: **<$900/kW for the complete SOFC system and <$225/kW for the stack**, with ≥40,000-hour stack life as a development goal.

I therefore use approximately $900/kW before fuel storage for the two mature SOFC systems. For PCFC I use $1,100/kW as a **development-target assumption**, because there is not yet a bankable large-marine PCFC cost basis. Ammonia tank size is adjusted for each architecture's efficiency.

For conventional propulsion I use the published $600/kW large marine engine assumption and HFO tank cost, add SCR/NO_x treatment, and add a scrubber for HSFO. A separate marine-fuel study uses $133/kW for SCR and cites SOFC costs of $925–1,315/kW. IMO's scrubber studies show a broad range of new-build exhaust-gas-cleaning-system costs; I use about $250/kW as a conservative 2026 planning allowance for the HSFO system.

That produces:

Propulsion system	Assumed installed propulsion + fuel-storage CAPEX	50-MW ship	100-MW ship
Direct SOFC + ammonia	~$1,228/kW	$61.4M	$122.8M
Partial-cracking SOFC + ammonia	~$1,223/kW	$61.2M	$122.3M
PCFC + ammonia	~$1,462/kW	$73.1M	$146.2M
HSFO diesel + tank + SCR + scrubber	~$1,220/kW	$61.0M	$122.0M
VLSFO diesel + tank + SCR	~$970/kW	$48.5M	$97.0M

Notice something surprising: **the mature partial-cracking SOFC/ammonia propulsion plant and the HSFO/scrubber propulsion plant come out at almost exactly the same initial capital cost in this model.**

That's not because fuel cells are intrinsically cheap today. It is because we're assuming the large-volume, mature SOFC cost target, while an HSFO ship needs the diesel, fuel tanks and substantial exhaust-treatment equipment.

Fuel-cell replacement is important, but not overwhelming

For SOFC stacks I modeled a conservative replacement schedule based on the marine SOFC paper: initial ammonia-SOFC life of about **5 years**, with each subsequent generation having 25% longer life and 42% lower replacement cost. I combine that schedule with DOE's mature **$225/kW stack-cost target**.

That produces SOFC replacements at approximately years 5, 11.3 and 19.1. Discounted over 25 years, the replacement burden adds only about:

> $7.5/MWh-shaft.

That may initially sound low. The reason is that we're replacing **the electrochemical stacks**, not buying a completely new $900/kW propulsion system every five years. The ammonia tanks, motors, power electronics, cracker pressure vessels, heat exchangers, structure and most balance-of-plant equipment remain aboard.

PCFC uncertainty is much greater. I used a deliberately higher $300/kW first stack replacement and obtained approximately **$10/MWh-shaft** for replacements. I would regard the $189/MWh PCFC total as much less certain than the SOFC numbers.

What 50 MW versus 100 MW means

Using the 31.6% full-power-equivalent operating profile, a 50-MW propulsion system delivers about **138,000 MWh of shaft energy annually**; a 100-MW system delivers about **277,000 MWh/year**.

Discounting 25 years of those outputs at 7%, the corresponding lifetime present-value expenditures are:

System	50-MW ship: 25-year PV cost	100-MW ship: 25-year PV cost
Direct SOFC + \$365/t NH_3	\$267M	\$535M
Partial-cracking SOFC + \$365/t NH_3	\$266M	\$532M
PCFC + \$365/t NH_3	\$304M	\$609M
HSFO + scrubber	\$247M	\$494M
VLSFO	\$295M	\$590M

The \$/MWh number is nearly identical at 50 and 100 MW in my calculation because I've conservatively scaled equipment approximately linearly with power. In reality, the 100-MW system should get some economies of scale in tanks, heat exchangers, controls and other balance-of-plant equipment, so the 100-MW ammonia case might be a few percent cheaper per MWh.

Why partial cracking still looks best

The direct-ammonia SOFC and partial-cracking SOFC are effectively tied economically: **\$166 versus \$165/MWh**. That difference is well below the uncertainty in the underlying assumptions.

I would nevertheless choose the **partial external cracker + SOFC**.

At ~600–700°C, hot SOFC exhaust can provide almost all the thermal duty for the ammonia cracker, ammonia vaporizer and preheater. Therefore, cracking does not require buying large amounts of additional fuel or electricity. Yet partial cracking reduces ammonia exposure inside the fuel-cell stack, gives much better control over thermal gradients and makes the catalyst something that can be serviced independently of the expensive electrochemical stack.

In other words, you give up very little efficiency or money in return for what I expect would be materially better durability and maintainability.

An especially useful break-even number

With the assumptions above, the optimized partial-SOFC ship costs about **\$165/MWh-shaft**.

At today's $588.50/t Singapore HSFO price, the HSFO ship is about $153/MWh. For the ammonia ship to exactly match it, green ammonia would need to fall from your assumed **$365/t to roughly $326/t**.

That's only about an **11% reduction in ammonia price**.

Conversely, holding ammonia at $365/t, HSFO would reach parity at roughly **$650/t**. That's not remotely an implausible bunker-fuel price historically.

Against VLSFO the situation is already reversed. The green-ammonia system breaks even with VLSFO at an ammonia price of roughly **$426/t**. Thus, at $365/t ammonia:

> Green NH_3 + partial SOFC: ~$165/MWh-shaft

versus:

> VLSFO: ~$183/MWh-shaft.

So, ammonia has roughly an **$18/MWh-shaft advantage**, or about 10%.

One important cost I have not included

There is a potentially meaningful penalty that does **not** show up in $/MWh-shaft: lost cargo space.

The published 56.38-MW ship analysis estimated roughly **32,760 m³ of ammonia tankage versus 15,142 m³ of HFO tankage**, while its SOFC machinery occupied about 3,132 m³ versus 1,289 m³ for an HFO engine. The authors explicitly identify the low volumetric power density of SOFC propulsion as a potential cargo-space problem.

Our 65%-efficient system would reduce ammonia tank volume somewhat compared with their 60% case, but ammonia still needs approximately twice the fuel-storage volume of HFO for the same voyage. If those extra tanks displace paying containers, the economically relevant metric becomes **$/TEU-mile**, not merely $/MWh-shaft.

Conversely, I have also **not imposed a carbon price on HSFO or VLSFO**. The same study uses an HFO combustion factor of 3.114 tonnes of CO_2 per tonne of fuel. A $100/t-$CO_2$ charge would add roughly **$55–60/MWh-shaft** to fossil propulsion, immediately making the $165/MWh ammonia/SOFC system considerably cheaper.

Bottom line

For the technology you are envisioning—**$365/t factory-produced nuclear green ammonia plus a highly optimized, waste-heat-integrated partial-cracking SOFC electric propulsion plant**—I would currently use **about $165/MWh of shaft energy** as the central 25-year planning number.

That compares with roughly **$153/MWh for new HSFO + scrubber** and **$183/MWh for new VLSFO** at today's Singapore bunker prices.

The result suggests something more consequential than just “green shipping costs more”: **if the $365/t automated-PWR ammonia target and ~$900/kW mature marine SOFC system can actually be achieved, zero-direct-CO_2 ammonia propulsion could plausibly be cheaper than VLSFO without any carbon subsidy and only modestly more expensive than HSFO.** The remaining economic question I would investigate next is the **cargo-space penalty in $/TEU-mile**, because that is probably now a larger economic uncertainty than the cracker itself.

Liquid Ammonia Production Quantity

Prompt to AI: Ammonia is currently used to make fertilizer for agriculture. Also, ammonia could be used to power large ships by feeding liquid ammonia into a high temperature fuel cell (HTFC) on ships, converting it to electricity, and then feeding that into an electrical motor. Also, ammonia could be used to power aerosol production on ships that implement Marine Cloud Brightening (MCB), to solve the global warming problem (i.e. create salt water aerosol, via electricity, via high temperature fuel cell, via ammonia). In units of millions of tonnes per year (Mt/year ammonia), what is the current: (a) global consumption of ammonia, (b) amount that would be needed to power large ships (via HTFC), and (c) amount that would be needed for MCB (via HTFC). If we replaced current global ammonia production with green ammonia (made without emitting CO2), then what percent of global CO2 emissions would be eliminated? If we replaced fossil fuel on large ships with green ammonia, then what percent of CO2 emissions would be eliminated?

Response From AI:

Page 94–96 of *Discussion, Electricity and Ammonia Cost* (google drive, AI Conversation) uses 1,500 TWh/year of electricity as the central Marine Cloud Brightening (MCB), with a very broad 400–3,000 TWh/year range. At the document's 55% net ammonia-to-electricity efficiency, 1,500 TWh corresponds to about 528 Mt of ammonia per year (i.e. “Mt” refers to millions of tonnes).

Ammonia use	Central requirement (million tons NH_3/year)	Basis
A. Current global ammonia consumption	~190 Mt/y	Current global production/consumption proxy
B. Power large ships with green NH_3 + optimized HTFC	~380 Mt/y	Ships ≥5,000 GT; ~65% ammonia-to-useful-power efficiency
C. MCB aerosol production, using document assumption	~528 Mt/y	1,500 TWh/y electricity; 55% HTFC efficiency
C-alt. MCB with our newer optimized 65% HTFC	~447 Mt/y	Same 1,500 TWh/y, but 65% efficiency
A + B + C	~1,100 Mt/y	Using the document's 55% MCB case
A + B + C-alt.	~1,015 Mt/y	Using 65% HTFC for both ships and MCB

So, MCB could be one of the largest ammonia uses in the world, comparable to or larger than the entire shipping requirement.

A. Current ammonia: about 190 Mt/year

The International Fertilizer Association estimates that global ammonia output reached **190.5 Mt in 2024**. At the global level, production is a reasonable proxy for consumption because inventories are small relative to annual production.

So, I would round this to:

> Current ammonia consumption ≈ 190 Mt NH_3/year.

Most of this ultimately supports nitrogen fertilizer production, although ammonia also has industrial uses.

B. Large ships: approximately 380 Mt NH_3/year

The latest IMO DCS data indicate that ships of **5,000 gross tonnes and above used about 223 Mt of fuel in 2024**. These ships account for more than 85% of international shipping emissions. The 2024 increase partly reflected longer voyages associated with Red Sea/Suez rerouting.

For a simple energy-equivalent conversion, I use approximately 40–41 MJ/kg for today's oil-dominated marine fuel mix, about 50% useful conversion efficiency for large modern marine engines, **18.6 MJ/kg LHV for ammonia**, and our optimized partial-cracking SOFC/HTFC system at about **65% ammonia-to-useful electrical/mechanical energy**.

That gives approximately:

> ≈ 380 Mt NH_3/year to replace the fossil fuel currently consumed by the large-ship fleet.

I would regard roughly **350–420 Mt/year** as a reasonable planning range because average existing-engine efficiency and future HTFC efficiency are uncertain.

C. MCB: approximately 528 Mt/year using your document

This is the major correction.

Your document says the central MCB planning requirement is:

> 1,500 TWh-electric/year

with approximately **400–3,000 TWh/year** as the preliminary range. It derives this from climate-model particle requirements and measured power demand from a first-generation aerosol nozzle, while emphasizing that spray-generator efficiency is highly uncertain.

The document then assumes 55% net ammonia-to-electricity efficiency. Ammonia contains approximately 5.17 MWh_{th}/t on an LHV basis. Therefore 1,500 TWh_e / (0.55 × 5.17 MWh_{th}/t) ≈ 528 Mt NH_3/year.

That is exactly the conclusion reached on pages 95–96 of the file: **about 530 Mt NH_3/year** for the central MCB spray-system case.

The document's complete MCB range is approximately 141 Mt/y at 400 TWh, 528 Mt/y at 1,500 TWh, and 1.06 billion tonnes/year at 3,000 TWh. It also notes that adding 15–25% for propulsion and auxiliaries would raise the central all-in numbers to roughly 607–660 Mt NH_3/year.

What if we use the newer 65%-efficient HTFC?

Since our subsequent ship analysis developed a more aggressive optimized high-temperature fuel-cell system at about 65%, it is useful to update the same 1,500-TWh MCB requirement.

Your existing document: ~528 Mt/year, using 55%.

Our newer optimized HTFC assumption: ~447 Mt/year, using 65%.

I would **not silently change the document's 530-Mt number**, because 55% is the explicit assumption behind it. But for future engineering work, I think **~450 Mt/year** is a reasonable aspirational central number if the high-efficiency SOFC system actually reaches 65% net electricity output.

The scale is much larger than I previously indicated

Using the original document assumptions: 190 + 380 + 528 ≈ 1,098 Mt NH_3/year,

or approximately 1.1 billion tonnes of ammonia per year.

That's about:

1,098 / 190 ≈ 5.8× today's entire ammonia industry.

Using the newer 65%-efficient HTFC for MCB: 190 + 380 + 447 ≈ 1,017 Mt NH_3/year,

or about 1.0 billion tonnes/year, approximately 5.3× present global ammonia production.

So, an appropriate high-level table is:

Application	NH_3 demand, Mt/y	Multiple of today's ~190 Mt/y production
Existing ammonia uses	190	1.0×

Large shipping	~380	2.0×
MCB, document 55% HTFC case	~528	2.8×
Combined	~1,100	~5.8×

This makes MCB very different from my previous answer. Under the assumptions on page 95, MCB alone would require almost three times today's entire world ammonia production.

How much global CO_2 would green ammonia for today's ammonia market eliminate?

The IEA estimates direct CO_2 emissions from ammonia production at about 450 Mt CO_2/year. It also identifies indirect emissions, including electricity used in ammonia production and downstream CO_2 release from urea application.

For the question "make today's ammonia without emitting CO_2," I would count the approximately 450 Mt of direct production emissions and roughly another 40 Mt/year associated with production electricity, but **not** downstream CO_2 from applying urea to fields.

That gives approximately:

> ~490 Mt CO_2/year avoided.

The Global Carbon Budget projects **38.1 Gt of fossil CO_2 emissions in 2025**.

So:

> Replacing today's fossil-based global ammonia production with essentially zero-CO_2 green ammonia would eliminate about 1.2–1.3% of global fossil CO_2 emissions.

That is consistent with the IEA's own characterization that ammonia production accounts for about **1.3% of energy-system CO_2 emissions**.

How much global CO_2 would green-ammonia shipping eliminate?

The latest large-ship fleet consumed about **223 Mt of fuel in 2024**. IMO's carbon factors are approximately 3.114 t CO_2/t HFO, 3.151 for light fuel oil, 3.206 for diesel/gas oil, and 2.750 for LNG.

Using the reported 2024 fuel mix gives an approximate large-ship fleet emission of:

> ~0.70 Gt CO_2/year.

Against 38.1 Gt/year of global fossil CO_2:

Therefore:

> Replacing the fossil fuel used by the world's large ships with genuinely green ammonia would eliminate approximately 1.8% of global fossil CO_2 emissions on a tank-to-wake basis.

The actual lifecycle reduction would be slightly smaller than 100% because nuclear/renewable ammonia production, transport and ship construction have some lifecycle emissions, even though ammonia itself contains no carbon.

Corrected overall picture

The numbers I would now carry forward are:

Item	Central value
Existing world NH_3 consumption	~190 Mt/year
NH_3 for large ships using optimized HTFC	~380 Mt/year
NH_3 for MCB using page-95 1,500 TWh and 55% HTFC	~528 Mt/year
MCB using 65% optimized HTFC instead	~447 Mt/year
Total NH_3, using document MCB assumption	~1.10 billion t/year
Today's ammonia decarbonization	~1.3% of global fossil CO_2 eliminated
Large-ship decarbonization	~1.8% eliminated
Both together	~3.1% of global fossil CO_2 eliminated

The central strategic insight is therefore quite different from my previous answer: shipping and MCB would each create ammonia markets substantially larger than today's fertilizer/industrial ammonia market. Under the page-95 MCB assumption, MCB is actually the largest of the three at roughly 530 Mt/year, followed by shipping at roughly 380 Mt/year, and today's ammonia industry at roughly 190 Mt/year.

And the 1,500 TWh/year MCB number itself remains highly uncertain. Your document explicitly treats 400–3,000 TWh/year as a preliminary range because aerosol-generator efficiency and the number of effective cloud-condensation nuclei required are unresolved. That uncertainty is much more consequential than whether the HTFC ultimately achieves 55% or 65% efficiency.

Purchase Orders

Prompt to AI: For the 11 technical work packages, list illustrative procurement packages, amounts, recipient types, and acceptance logic. Treat "purchase order" broadly as a contract, grant, cooperative agreement, standards-development agreement, reimbursable agreement, or task order.

Response From AI:

Every award should define a controlled deliverable, interfaces, test evidence, data rights, schedule, payment milestones, safety/quality requirements, and an explicit stop/advance rule. The following packages reconcile to the $800M planned technical program; the $200M reserve remains sponsor-controlled and is not pre-awarded.

PO	Work package	Amount	Illustrative recipient type	Acceptance logic
1-1	Reference architecture, 31-platform portfolio, and digital twin	$22M	Lead systems integrator / owner engineer	Controlled architecture, 31-platform requirements and digital model; accept after independent requirements review.
1-2	Reference architecture, 31-platform portfolio, and digital twin	$32M	Competing platform/system design teams	FEED, interface definitions and economic screening; accept against design reviews and reproducible models.
1-3	Reference architecture, 31-platform portfolio, and digital twin	$16M	Independent reviewer / cost team	Architecture, hazard and economics replication; reports directly to owner.
2-1	31 industrial-processing platform designs and pilot prototypes	$45M	Process/EPC/university design teams	Thirty-one controlled platform designs; first-five detailed designs; accept after PDR/CDR and interface conformance.
2-2	31 industrial-processing platform designs and pilot prototypes	$110M	Platform fabricators / OEMs / research teams	Functional pilots and endurance campaigns; payment tied to witnessed operating data and teardown evidence.
2-3	31 industrial-processing platform designs and pilot prototypes	$25M	Independent test laboratories	Replicate efficiency, throughput, degradation, safety and maintainability results.
3-1	Hydrogen and ammonia production platforms	$20M	Electrolyzer/HTSE and heat-integration team	SOEC/steam integration and oxygen/water systems; accept on efficiency/degradation tests.
3-2	Hydrogen and ammonia production platforms	$55M	Ammonia licensors / ASU / compressor / storage consortium	Integrated hydrogen-to-liquid-ammonia pilot; accept on purity, availability, safety and repeated production.

PO	Work package	Amount	Illustrative recipient type	Acceptance logic
3-3	Hydrogen and ammonia production platforms	$20M	Independent chemical/process safety team	Cost replication, emissions/water accounting, HAZOP and long-duration inspection.
4-1	Marine ammonia fuel-cell propulsion prototypes and endurance testing	$15M	SOFC/PCFC/cracker competing teams	Three propulsion architectures and durability rigs; accept on agreed load cycles and emissions.
4-2	Marine ammonia fuel-cell propulsion prototypes and endurance testing	$40M	Marine integrator + fuel-cell OEM + motor/inverter supplier	100–500 kW prototypes and MW-class shared demonstrator; accept on net efficiency and stable operation.
4-3	Marine ammonia fuel-cell propulsion prototypes and endurance testing	$15M	Classification society / independent marine lab	Witness endurance, teardown, emissions, safety and marine integration tests.
5-1	Material handling, tunnels, storage, waterfront, and modular logistics equipment	$20M	Industrial logistics/EPC design team	Distribution spine, tunnels, storage, waterfront and modular equipment designs.
5-2	Material handling, tunnels, storage, waterfront, and modular logistics equipment	$50M	Fabricators / port equipment / conveyor / piping suppliers	Representative modules and integrated material-flow prototypes.
5-3	Material handling, tunnels, storage, waterfront, and modular logistics equipment	$15M	Independent logistics/safety reviewer	Throughput, damage, contamination, maintenance and recovery audits.
6-1	Heavy module transport, double rail, ship/shore transfer, and gantry systems	$15M	Heavy transport / structural / rail engineering team	Rigid truss, carrier, gantry, route and ship-interface design.
6-2	Heavy module transport, double rail, ship/shore transfer, and gantry systems	$40M	Heavy transport fabricator / shipyard / rail supplier	Scaled and representative loaded transfer systems.
6-3	Heavy module transport, double rail, ship/shore transfer, and gantry systems	$10M	Marine warranty / structural independent reviewer	Witness proof loads, curves/grades, ship transfer and recovery cases.
7-1	Electrified rail and high-power construction-machine interface	$10M	Electrical/rail systems design consortium	Voltage classes, enclosed conductor, segment logic, safety and interoperability design.
7-2	Electrified rail and high-power construction-machine interface	$22M	Switchgear / conductor / collector / controls suppliers	Full-voltage test loop and mobile collectors; fault and regenerative tests.
7-3	Electrified rail and high-power construction-machine interface	$8M	Independent HV safety lab / standards body	Grounding, water/dirt, fault clearing, lockout and conformance verification.
8-1	Automated construction of	$10M	Tank/EPC/robotics design	Jacking, panel handling, welding/NDE

PO	Work package	Amount	Illustrative recipient type	Acceptance logic
	thermal-storage tanks		consortium	and sequence design.
8-2	Automated construction of thermal-storage tanks	$32M	Robotics, welding, NDE and heavy-lift fabricators	Representative tank and machine prototypes with full-size weld trials.
8-3	Automated construction of thermal-storage tanks	$8M	Independent tank/code/NDE reviewer	Witness weld/NDE, dimensional, jacking and production-rate tests.
9-1	Standards, interoperability laboratory, digital thread, and conformance testing	$15M	Standards secretariat + systems engineering team	Master architecture and open interface documents through v1.0.
9-2	Standards, interoperability laboratory, digital thread, and conformance testing	$28M	Interoperability lab + multi-vendor participants	Conformance fixtures, reference implementations and plug tests.
9-3	Standards, interoperability laboratory, digital thread, and conformance testing	$12M	ISO/IEC/ANSI/SAC/industry coordination and independent review	Consensus packages, publication, IP/data rules and v2.0 release.
10-1	Integrated prototype campus, flow/corrosion loops, and long-duration campaigns	$20M	Test-campus EPC / utilities team	Externally heated PWR/HTR loops, test bays, safety systems and instrumentation.
10-2	Integrated prototype campus, flow/corrosion loops, and long-duration campaigns	$30M	Materials/corrosion/test operations team	Representative flow loops, coupons/spools, long-duration operation and inspections.
10-3	Integrated prototype campus, flow/corrosion loops, and long-duration campaigns	$15M	Independent test/data team	Calibration, test witnessing, data quality, teardown and reliability analysis.
11-1	Program management, safety, techno-economics, and independent review	$8M	Managing organization / program office	Procurement, baseline, configuration, risk, finance and reporting.
11-2	Program management, safety, techno-economics, and independent review	$7M	Safety, environmental, cybersecurity and marine assurance teams	Cross-program hazard oversight and independent stop-work authority.
11-3	Program management, safety, techno-economics, and independent review	$10M	Independent economics/audit/review teams	Replicated cost/LCA models, annual technical audits and Phase-II decision report.

Initiative-specific contracting additions

- **Classify every major award and cost as platform-family-specific, shared industrial infrastructure, hydrogen/ammonia, marine propulsion, or shared with the Nuclear Power Automated Site-Construction initiative so shared-asset economics can be audited without double funding.**
- **Flow down the applicable process-safety, pressure-system, hazardous-area, ammonia/hydrogen, marine, electrified-rail, tank-construction, cybersecurity, environmental, and interoperability requirements in addition to the common portfolio terms.**
- **Acceptance packages for process and infrastructure prototypes must include the final configuration, native design/source files, raw test data, calibration, efficiency and throughput, degradation/corrosion evidence, failures/nonconformances, maintenance evidence, and a reproducibility package.**
- **Where the separate nuclear-construction initiative has already qualified a transporter, gantry, rail geometry, digital platform, metrology system, or other shared asset, this initiative funds only incremental industrial requirements, adaptation, operating time, or additional tests.**

All other procurement, payment, open-source/data, conflict, audit, change-control, reserve, shared-asset, and transition provisions are specified in the Moonshot Common Framework.

Purchase Order Recipients

Prompt to AI: Who might receive the fund's purchase orders? Identify illustrative organization types and examples spanning dual-reactor interfaces, the 31 process platforms, ammonia, marine fuel cells, giant-module logistics, electrified rail, thermal-storage construction, standards, and independent review.

Response From AI:

The names below are illustrative candidates, not endorsements or prequalified bidders. Competition should emphasize teams with real manufacturing/test capability, open interfaces, and willingness to provide auditable data. Independent reviewers should not be financially dependent on the hardware they evaluate.

Capability	Illustrative candidates	Potential role
Dual-reactor reference architecture and nuclear/industrial heat	Tsinghua INET; Chinese nuclear owner-engineers and design institutes; INL; ORNL; Argonne; major EPC firms	PWR/HTR utility envelopes, isolation heat systems, licensing/safety boundaries, systems engineering.
31 industrial processing platforms	Major Chinese and international EPC/process firms; heat-pump, MVR, dryer, distillation, alumina, furnace and industrial-equipment OEMs; universities	FEED, detailed platform designs, pilots, factory packages, thermal integration and process optimization.
Hydrogen and ammonia	SOEC/HTSE developers; Topsoe; KBR; Casale; thyssenkrupp Uhde/nucera; Air Liquide; Linde; compressor/refrigeration/storage suppliers	Electrolysis, nitrogen, synthesis, heat recovery, storage/loading and integrated ammonia pilot.
Marine ammonia fuel cells	SOFC/PCFC developers; marine electrical integrators; major shipyards; propulsion motor/inverter suppliers; catalyst/cracker firms	Direct/partial-cracking SOFC and PCFC prototypes, heat recovery, electric propulsion and durability.
Marine classification and safety	DNV; ABS; Lloyd's Register; Bureau Veritas; marine warranty surveyors	Safety cases, class interpretation, ammonia handling, fuel-cell installation, test witnessing.
Heavy transport and giant modules	Mammoet; Sarens; Goldhofer; TII Scheuerle; rail/gantry/heavy-civil contractors; shipyards	Rigid truss, synchronized carriers, double rail, jacking, gantries, ship/shore transfer.
Material handling and waterfront	Port engineering firms; conveyor/unloader manufacturers; pump/compressor firms;	Tunnels, manifolds, conveyors, transfer towers, ship loaders/unloaders, storage and

Capability	Illustrative candidates	Potential role
	tank/terminal EPCs	distribution spine.
Electrified rail and high-power machine power	Siemens, ABB, Schneider Electric, Hitachi Energy and Chinese switchgear/rail-electrification suppliers; IEC/rail labs	Enclosed conductors, LV/MV switchgear, collectors, grounding, segmentation, regenerative operation and safety.
Thermal-storage tank automation	Large tank EPCs; welding/NDE robotics suppliers; jack/lift manufacturers; ASME/API specialists	Tank jacking, automated welding/NDE, panel handling, foundation/floor/roof/insulation systems.
Standards and interoperability	ISO/IEC/ANSI/SAC/SAMR committees; ASME; IEEE; ISA; OPC Foundation; AISC; API; OCIMF	Interface standards, conformance tests, digital schemas, cyber/functional safety and publication.
Independent economics/testing	INL, NREL, NETL, Argonne, universities, independent cost engineers, materials labs and safety organizations	Replicate performance, corrosion/degradation, cost, lifecycle and safety results; report directly to owner.

Managing organization versus recipient

Managing Organization self-award, related-party, competition, pricing, conflict, audit, and recipient-transition rules are specified in the Moonshot Common Framework. For this initiative, ownership and affiliation among reactor/interface providers, process licensors, EPCs, equipment OEMs, ship and port firms, standards organizations, universities, test facilities, and Managing Organization entities must be visible so cost, interoperability, and technology comparisons remain credible.

Project Management

Prompt to AI: What project-management provisions are specific to this nuclear-integrated industrial-processing initiative after incorporating the Climate Moonshot Portfolio Common Framework?

Response From AI:

Portfolio-wide governance, Managing Organization selection, fiduciary controls, procurement, financial controls, reserve governance, independent review, audit, conflicts, open-source/data rights, shared-asset rules, and transition provisions are specified in the Moonshot Common Framework. This chapter states only the additional technical-management structure required by the industrial-processing program.

Entity	Initiative-specific responsibility	Program-specific limit / interface
Program Manager / PMO	Own one integrated five-year execution baseline across the 31-platform portfolio, eleven baseline work packages, shared infrastructure, and the non-nuclear prototype campus.	Uses the common-framework fiduciary and procurement controls; may not waive nuclear-interface boundaries, independent review, process-safety hold points, or reserved Funding Source decisions.
Chief Engineer / Industrial Platform Design Authority	Own the 12 m × 96 m platform architecture, common dimensions and utility interfaces, heat-cascade logic, systems engineering, configuration control, verification strategy, and standards baseline.	May optimize industrial equipment and common interfaces but cannot alter controlled reactor safety functions or licensed nuclear systems.
Nuclear / Industrial Heat Interface Authority	Own the PWR- and HTR-PM-side utility assumptions, isolation boundary, heat grades, metering, interface loads, and assumptions used for nuclear-to-industrial integration.	Reactor-side changes require the applicable nuclear design authority and regulator process; the five-year prototype program remains non-nuclear unless separately authorized.
Process Safety / Chemical Operations Authority	Own PHA/HAZOP, pressure relief, hazardous areas, hydrogen/ammonia inventories, fire protection, operating procedures, emergency response, and chemical test readiness.	Has independent hold/stop authority for hazardous process operation and cannot be overruled solely for schedule or cost reasons.
Infrastructure / Marine / Heavy-Logistics Leads	Integrate tunnels, material handling, 12×96 m transport, waterfront transfer, electrified rail, thermal-storage-tank construction, and site choreography.	Must coordinate shared assets and interfaces with the nuclear-construction initiative and document when separate industrial equipment is actually justified.
Independent Test / Standards / Techno-Economics Authorities	Witness or independently reconstruct performance, durability, corrosion, emissions, interoperability, transport, safety, cost, and climate-benefit evidence.	Must remain institutionally separate from the performer whose claim is accepted and retain direct access to raw/source evidence.

Entity	Initiative-specific responsibility	Program-specific limit / interface
Prototype Campus / Phase II Transition Leads	Run the non-nuclear PWR-like and HTR-like heat loops, integrated campaigns, maintenance and endurance testing, and prepare platform-family transition packages.	Phase II, a nuclear-connected demonstration, commercial chemical production, or ship deployment is separately authorized; evidence may justify scaling some platform families and stopping others.

Initiative-specific management requirements

- The Managing Organization must be capable of large industrial R&D procurement and integration spanning process equipment, hydrogen/ammonia, marine systems, heavy logistics, standards, long-duration testing, and China/international supply chains.
- Maintain one controlled 31-platform reference architecture and interface-control baseline covering module dimensions, heat, piping, electricity, rail, controls/data, metrology, storage, waterfront transfer, and safety boundaries.
- Independent acceptance is required for high-temperature/pressure loops, ammonia and fuel-cell endurance, corrosion/materials evidence, giant-module handling, electrified-rail high-power tests, tank welding/NDE, multi-vendor interoperability, and bottom-up cost/emissions models.
- Use the common framework's one-funded-asset rule for equipment shared with the Nuclear Power Automated Site-Construction initiative; a joint interface/change board should control shared geometry, loads, digital identifiers, and test requirements.
- By Year 4-5, maintain separate Phase II packages for the platform families and enabling systems that survive technical, safety, cost, durability, and market screening. Follow-on funding need not be distributed evenly across the original 31 designs.

Funding Source Requirements

We assume the Funding Source aims to maximize climate benefit per dollar spent, as opposed to maximize financial return on an investment. This is sometimes referred to as “climate money”, and it differs from the traditional investment portfolio. Sources of climate money include high-net-worth individuals, foundations, and governments.

The portfolio-wide Funding Source requirements are specified in the Moonshot Common Framework. The additions below address the 31-platform industrial mission, dual-reactor planning assumptions, chemical and marine hazards, open interoperability, long-duration testing, and the shared-asset boundary with the nuclear-construction initiative.

Requirement	Initiative-specific addition to the common framework
Industrial-processing mission	Spend only on the five-year R&D mission for the 31 industrial-processing platform families and the supporting ammonia, marine, logistics, rail, thermal-storage, standards, and non-nuclear test infrastructure defined by this proposal.
Dual-reactor planning scenarios	Treat Automated China PWR and HTR-PM electricity/heat values as explicit planning scenarios to be tested, not guaranteed commercial prices or authorization to build a reactor.
Nuclear / industrial isolation	Maintain a qualified isolation boundary so process, chemical, electrical, logistics, or platform faults cannot compromise reactor safety. Any nuclear-connected demonstration requires separate licensing and authorization.
Evidence-based platform down-selection	Do not require all 31 platform families to receive equal prototype scale or follow-on funding. Increase spending only where measured efficiency, durability, safety, economics, and climate value justify it.
Process and hazardous-energy safety	Apply independent PHA/HAZOP, relief/containment, hazardous-area, fire/explosion, ammonia/hydrogen, pressure/high-temperature, emergency-response, worker-safety, and environmental controls appropriate to each test.
Marine ammonia and waterfront work	Use marine-class, bunkering, fuel-system, emissions, port, ship/shore-transfer, and warranty/safety review appropriate to the actual prototype; land-based success does not by itself establish ship readiness.
Open standards / lawful process-IP exceptions	Keep sponsor-funded interface standards, conformance tools, digital schemas, and reference implementations open and multi-vendor. Identify process-licensor and other background-IP restrictions before award and narrow them to what is genuinely pre-existing.
Shared assets / no double funding	Apply asset-level accounting to transporters, gantries, rail, digital systems, metrology, test facilities, and other capabilities shared with the nuclear-construction initiative; the second program pays only incremental requirements, adaptation, operating time, or consumables.
Long-duration and degradation evidence	Require representative temperature, pressure, chemistry, duty cycle, corrosion, catalyst/stack aging, sealing, cycling, maintenance, and teardown evidence before

	scaling systems whose economics depend on long life.
Transparent product economics	Publish consistent accounting boundaries for electricity, useful heat, water, hydrogen, ammonia, DAC, shipping energy, storage, logistics, capital, replacement, financing, emissions, and product delivery so nuclear and fossil comparisons can be independently reconstructed.
Global transferability and policy accounting	Maintain selected U.S. and European deployment sensitivities using the same measured platform evidence. Re-price local reactor/heat interfaces, energy and feedstocks, platform localization, codes/qualification, labor, transport, financing, carbon costs, subsidies, tariffs, and product-delivery boundaries; report engineering economics separately from policy-adjusted economics.
Separate Phase II authorization	Completion of the five-year R&D program does not authorize commercial chemical production, a nuclear-connected industrial yard, a commercial ammonia ship fleet, or equal follow-on funding for every platform family.

Program-specific record additions

- Controlled 31-platform requirements, priority order, interface definitions, heat/electricity assumptions, conventional fossil baselines, and change history.
- As-built/as-run process, ammonia, fuel-cell, logistics, rail, tank, corrosion, and prototype-campus configurations; calibration; failures; nonconformances; maintenance; teardown findings; and corrective actions.
- An indexed record of nuclear-interface assumptions, process-safety analyses, marine/classification assumptions, environmental permits, cybersecurity controls, standards/conformance results, and any lawful controlled-data or background-IP restrictions that affect reuse.
- Side-by-side estimates for delivered product cost and CO2 reduction across relevant PWR, HTR-PM, coal, natural-gas, and other incumbent cases, with consistent financing, capacity factor, product specification, infrastructure, and delivery boundaries. Include selected U.S. and European transferability sensitivities showing local reactor/heat-source basis, labor, platform localization, codes/qualification, financing, policy, logistics, and feedstock effects.

Climate-benefit metric

Do not claim climate benefit from prototype completion alone. Report the expected reduction in delivered cost and deployment time of low-carbon hydrogen or ammonia, multiplied by plausible deployment scale and adjusted for technical, licensing, financing, infrastructure, market, and adoption risk. Publish ranges and assumptions. Also report whether the platform system reduces or increases the number of unique site designs, construction labor, outage exposure, water use, hazardous inventories, and product-transport distance.

Cost-model equations and bookkeeping

- Hydrogen: levelized cost per kilogram equals annualized capital, fixed O&M, variable O&M, electricity, heat, water, stack replacement, compression, storage, transport, financing, and decommissioning, minus defensible co-product credits, divided by annual saleable hydrogen.
- Ammonia: levelized cost per metric ton adds hydrogen, nitrogen, synthesis-loop, refrigeration, catalyst, storage, loading, losses, and financing. Stoichiometrically, one metric ton of ammonia contains about 176 kg of hydrogen and 824 kg of nitrogen before process losses.
- Water and oxygen: splitting water requires at least 9 kg of water per kg of hydrogen and produces 8 kg of oxygen per kg of hydrogen. Actual water use is higher because of purification, cooling, blowdown, and losses.
- Comparisons must use the same financing year, currency, capacity factor, product pressure/purity, plant life, tax, infrastructure boundary, and delivery point.

Top Industrial Processing and Nuclear Integration References

Prompt to AI: Identify and summarize the papers, reports, standards, or major references most relevant to the fund. Explain why each matters and what it does not establish.

Response From AI:

The list emphasizes references most useful for nuclear-integrated industrial processing: reactor/process-heat interfaces, heat pumps and electrification, hydrogen/ammonia, DAC, marine propulsion, heat exchangers, thermal storage, interoperability, and heavy modular transport. It is a program-design list rather than a citation-count ranking.

Reactor and Nuclear Process-Heat Foundations

Zhang, "Design Aspects of the Chinese Modular High-Temperature Gas-Cooled Reactor HTR-PM," 2006

Establishes the design lineage, module arrangement, helium-cooled reactor concept, and safety philosophy relevant to defining the nuclear utility envelope. Limitation: Orientation only; not a design-data license, process-heat interface standard, or commercial platform design. | Link

Zhang, "Current Status and Technical Description of Chinese 2 × 250 MWth HTR-PM Demonstration Plant," 2009

Provides technical description of the HTR-PM demonstration plant and useful temperatures, power levels, systems, and module relationships for reference architecture. Limitation: Published information cannot replace controlled vendor drawings, operating limits, or regulator-approved interface data. | Link

Zhang, "The Shidao Bay 200 MWe HTR-PM Demonstration Power Plant: An Engineering and Technological Innovation," 2016

Explains the realized plant and helps separate proven reactor features from proposed industrial-platform additions. Limitation: Does not establish cost or licensing of a replicated plant outside China. | Link

Zhang, "Loss-of-Cooling Tests Validate the Inherent Safety of the World's First Commercial Small Modular Reactor," Joule, 2024

Reports full-scale loss-of-cooling test evidence important to the user's choice of an HTR-PM reference. Limitation: Reactor safety evidence does not make a co-located hydrogen or ammonia facility safe; chemical hazards and coupling must be analyzed separately. | Link

IAEA, "Opportunities for Cogeneration with Nuclear Energy," 2017

Frames nuclear cogeneration for process heat, hydrogen, desalination, and industrial applications and emphasizes matching reactor temperatures to process needs. Limitation: Does not prescribe the proposed open platform interface or transport system. | Link

Arvanitidis, "Nuclear-Driven Integrated Energy Systems: A State-of-the-Art Review," 2023

Reviews nuclear integration with hydrogen, ammonia, fuels, thermal storage, grids, and industrial processes, helping define interfaces and research gaps. Limitation: A broad review cannot replace detailed design of one reactor, product, platform, route, and site. | Link

Ramirez-Corredores, "Identification of Opportunities for Integrating Chemical Processes for Carbon (Dioxide) Utilization to Nuclear Power Plants," 2021

Examines process integration, temperature matching, heat and electricity use, and chemical opportunities adjacent to nuclear plants. Limitation: Focuses on process opportunities rather than the proposed transport and open-interface hardware. | Link

Bragg-Sitton, "Non-Electric Applications of Generation IV Reactors," 2022

Provides a systematic view of reactor temperature ranges, industrial heat demands, augmentation, and non-electric products. Limitation: Does not prove economic deployment or settle the reactor-to-process licensing boundary. | Link

Industrial Heat Pumps and Process Electrification

IEA, "The Future of Heat Pumps in China," 2024

Directly supports the program's emphasis on low- and medium-temperature industrial heat: the report identifies major Chinese industrial demand below 200°C, highlights heat-pump potential in light industry, pulp and paper, and chemicals, and notes standardization as a barrier that the platform concept is intended to address. Limitation: The report evaluates heat-pump deployment and policy at sector scale; it does not establish the cost, COP, reliability, or integration performance of a particular nuclear-powered 12 m × 96 m platform. | Link

Hydrogen and High-Temperature Electrolysis

IAEA, "Hydrogen Production Using Nuclear Energy," Nuclear Energy Series, 2013

Comprehensive reference on nuclear hydrogen pathways, reactor coupling, high-temperature electrolysis, thermochemical cycles, safety, economics, and deployment considerations. Limitation: Technology and cost assumptions require updating and site-specific validation. | Link

O'Brien, "High-Temperature Electrolysis for Large-Scale Hydrogen and Syngas Production from Nuclear Energy," 2009

A key technical treatment of coupling high-temperature gas-cooled reactors to solid-oxide electrolysis and using heat to reduce electrical demand. Limitation: Large commercial stack lifetime, cost, manufacturing yield, and availability remain separate questions. | Link

O'Brien, "High Temperature Electrolysis for Hydrogen Production from Nuclear Energy – Technology Summary," INL, 2010

Summarizes thermodynamics, cell technology, system efficiency, test results, and reference integration for nuclear-powered HTSE. Limitation: A technology summary is not a bankable vendor guarantee or a current commercial quote. | Link

O'Brien, "High Temperature Steam Electrolysis: Demonstration of Advanced Solid-Oxide Stack Technology," INL, 2011

Documents laboratory-scale integrated operation and degradation behavior, informing pilot design and acceptance metrics. Limitation: Scale and duration are far below a 100+ MW commercial platform. | Link

Harvego, "Economic Analysis of the Reference Design for a Nuclear-Driven High-Temperature-Electrolysis Hydrogen Production Plant," INL, 2008

Provides a detailed reference flowsheet and H2A-based economic method for a high-temperature helium reactor coupled to electrolysis. Limitation: Its dollar values, financing, reactor costs, stack assumptions, and scale are dated; reuse the method, not the headline number. | Link

Frick, "Technoeconomic Assessment of Hydrogen Cogeneration via High Temperature Steam Electrolysis with a Light-Water Reactor," Applied Energy, 2022

Shows how thermal and electrical integration, market operation, capital cost, and plant utilization affect nuclear hydrogen economics. Limitation: A light-water-reactor case does not directly establish HTR-PM performance or platform logistics. | Link

Cheng, "Assessing the Role of High Temperature Steam Electrolysis in Nuclear-Integrated Hydrogen Production," 2026

Recent techno-economic assessment useful for current parameter ranges, sensitivity methods, and comparison of HTSE with lower-temperature electrolysis. Limitation: Results depend on assumed electricity opportunity cost, stack performance, financing, and utilization. | Link

U.S. DOE, "H2A Production Models" and methodology documentation

Provides standardized discounted-cash-flow definitions and inputs for levelized hydrogen cost, enabling transparent comparison and sensitivity analysis. Limitation: The model is only as credible as its inputs and system boundary; it does not independently validate hardware performance. | Link

U.S. DOE H2NEW Consortium, technical targets and durability research

Provides current performance, manufacturing, durability, and cost targets for electrolyzers and a framework for stack and system validation. Limitation: Targets are not achieved commercial performance and are not specific to HTSE or nuclear steam integration. | Link

Direct Air Capture

Agbo, "Technological Innovation Opportunities for CO2 Removal," U.S. DOE, 2024

Provides a current national-laboratory assessment of direct-air-capture cost drivers and R&D opportunities, including regeneration energy, sorbent lifetime and cost, contactor capital, capture kinetics, lifecycle accounting, and durable storage. It is a strong basis for deciding what the nuclear-powered DAC platforms should measure rather than assuming a single mature cost. Limitation: The cost-reduction cases are R&D analyses and illustrative scenarios, not guarantees that a particular DAC technology will reach $100/tCO2 or that nuclear integration alone will deliver the modeled reductions. | Link

Ammonia and Low-Carbon Fuels

Nørskov, "Sustainable Ammonia Synthesis," U.S. DOE Basic Research Needs Roundtable Report, 2016

Defines the scientific opportunity space for lower-energy ammonia synthesis, catalysts, separations, electrochemical routes, and conventional Haber-Bosch improvement. Limitation: Research opportunities are not commercial readiness claims; the five-year fund should use proven synthesis where needed and test alternatives separately. | Link

Smith, "Current and Future Role of Haber–Bosch Ammonia in a Carbon-Free Energy Landscape," Energy & Environmental Science, 2020

Explains ammonia's energy and fertilizer role, conventional process energy, low-carbon hydrogen integration, storage, and system tradeoffs. Limitation: Does not provide a nuclear-specific platform design or full logistics cost. | Link

MacFarlane, "A Roadmap to the Ammonia Economy," Joule, 2020

Broad roadmap covering production, use, transport, safety, direct synthesis, fuel applications, and research priorities. Limitation: Roadmap scenarios are not a substitute for product offtake contracts, local regulations, or platform-scale hazard analysis. | Link

Armijo, "Flexible Production of Green Hydrogen and Ammonia from Variable Solar and Wind Energy: Case Study of Chile and Argentina," 2020

Useful benchmark for how utilization, storage, electrolysis, air separation, and ammonia-loop flexibility affect low-carbon ammonia economics. Limitation: Variable-renewable integration differs from high-capacity-factor nuclear heat and electricity, so assumptions must be normalized. | Link

Marine Ammonia and Fuel-Cell Propulsion

Kanchiralla, "How Do Variations in Ship Operation Impact the Techno-Economic Feasibility and Environmental Performance of Fossil-Free Fuels? A Life Cycle Study," Applied Energy, 2023

Compares alternative marine fuels and propulsion choices under realistic variations in ship operation and is especially useful for testing whether ammonia-SOFC conclusions remain robust when utilization, voyage profile, lifecycle emissions, and system assumptions change. Limitation: It is a scenario and lifecycle study rather than a full-scale ammonia-SOFC sea trial, and its fuel prices, equipment costs, efficiencies, and operating profiles must be replaced with project-specific values for a 50–100 MW ship. | Link

Micoli, "Advancing Zero-Emission Maritime Solutions: Case Study of an Ammonia-Powered Fuel Cell System Implementation," Energy Conversion and Management: X, 2024

Provides a ship-level implementation case for ammonia-derived fuel-cell power, including the practical consequences of fuel processing, machinery arrangement, storage volume, weight, and ship integration. It helps connect fuel-cell efficiency calculations to physical vessel design. Limitation: A case study on one vessel class does not establish MW-scale commercial durability, class approval, stack replacement cost, or the economics of the specific high-temperature SOFC/PCFC architectures proposed in this fund. | Link

Heat Exchangers, Nuclear Cost, and Modular Integration

ASME, "Boiler and Pressure Vessel Code Roadmap for Compact Heat Exchangers in Advanced Reactors," 2020

Identifies code, materials, fabrication, inspection, and in-service gaps for compact heat exchangers that may isolate nuclear heat from industrial loops. Limitation: A roadmap does not qualify a particular exchanger or create a complete interface standard. | Link

OECD Nuclear Energy Agency, "Reducing the Costs of Nuclear Power," 2020

Reviews standardization, repeat construction, modularization, project governance, financing, supply chains, and regulatory effects on nuclear cost. Limitation: General cost lessons do not establish the economics of HTR-PM replication or chemical platforms. | Link

Thermal Storage and Industrial Interoperability Standards

ASME TES-1-2023, "Safety Standard for Thermal Energy Storage Systems: Molten Salt"

Defines lifecycle requirements for molten-salt thermal-energy-storage systems, including storage vessels, pumps, valves, piping, heat exchangers, controls, inspection, testing, commissioning, operation, maintenance, and decommissioning. It gives the tank-automation work package an established safety framework to build upon rather than replace. Limitation: TES-1 does not specify the proposed robotic tank-jacking, panel placement, automated welding/NDE machines, modular site interfaces, or the economics of automated field construction. | Link

IEC 62541-1:2025, "OPC Unified Architecture - Part 1: Overview and Concepts"

Provides the foundation for OPC UA, the industrial interoperability architecture proposed in this document for common equipment identity, status, commands, alarms, maintenance, and data exchange across independently supplied processing platforms and site infrastructure. Limitation: The base OPC UA standard does not define this program's platform-specific semantic model, mechanical or thermal interfaces, conformance tests, cybersecurity zoning, or safety logic; those require a dedicated companion specification and integration standard. | Link

Heavy Module and Marine Transport

DNV-ST-N001, "Marine Operations and Marine Warranty," current edition

Provides a recognized framework for planning, design, loads, weather criteria, documentation, and verification of heavy marine operations. Limitation: Class and warranty standards must be applied to the actual cargo, vessel, ports, routes, and operations; they do not certify the proposed concept in advance. | Link

IMO, "Code of Safe Practice for Cargo Stowage and Securing" and Cargo Securing Manual guidance

Defines international expectations for cargo loading, stowage, securing, manuals, and shipboard responsibilities, including heavy and semi-standardized cargo. Limitation: Compliance does not solve the engineering of a 12 m by 96 m platform or guarantee port acceptance. | Link

Top Researchers and Organizations

Prompt to AI: Who are leading or especially relevant researchers, organizations, vendors, laboratories, standards bodies, transport firms, and process licensors? Group them by capability and explain their potential role.

Response From AI:

No single organization covers the mission. The fund needs HTR-PM and HTGR experts, nuclear owner engineers, process-heat and heat-exchanger specialists, electrolyzer researchers and vendors, ammonia licensors, modular fabrication yards, heavy transport and maritime firms, standards bodies, chemical and nuclear safety experts, cost modelers, regulators, and megaproject managers. “Top” below means relevant to the proposed work, not an objective global ranking. Current affiliation, availability, export restrictions, and conflicts must be verified before appointment or award.

Capability	Illustrative people / organizations	Potential contribution
HTR-PM / HTGR and process heat	Zhang Zuoyi and HTR-PM teams at Tsinghua INET; experts at IAEA HTGR knowledge programs; JAEA HTGR/HTTR teams; INL advanced-reactor and integrated-energy researchers.	Reference reactor envelope, high-temperature helium systems, safety, intermediate heat loops, and operating transients.
Nuclear integrated energy systems	Shannon Bragg-Sitton, Richard Boardman, Konor Frick, Daniel Wendt, Paul Talbot, J. E. O’Brien, and related INL/NREL/DOE teams.	Systems integration, HTSE, dynamic models, test facilities, cost analysis, process heat, and owner-side technical review.
Electrolysis and hydrogen	H2NEW consortium researchers; NREL electrolysis and H2A teams; solid-oxide electrolysis groups at INL, national laboratories, universities, and vendors such as Bloom, Sunfire, FuelCell Energy, Topsoe, and Siemens Energy.	Stack durability, manufacturing, steam integration, power electronics, balance of plant, compression, purity, and cost.
Ammonia science and process engineering	Jens Nørskov, Ib Chorkendorff, Douglas MacFarlane, Robert Schlögl, William Schneider, and conventional ammonia process experts at Topsoe, KBR, Casale, thyssenkrupp Uhde, Air Liquide, and Linde.	Catalysts, Haber-Bosch design, flexible operation, air separation, heat integration, refrigeration, storage, and alternative synthesis review.
Heavy transport and marine operations	Senior engineers at Mammoet, Sarens, Goldhofer, TII Scheuerle, DNV, ABS, Lloyd’s Register, Bureau Veritas, heavy-lift shipping companies, port engineers, and major shipyards.	Carrier dynamics, route engineering, active jacking, dock transfer, vessel stability, seafastening, weather criteria, class, and recovery.
Modular construction and	Module fabrication and construction leaders	Design-for-manufacture-and-assembly,

Capability	Illustrative people / organizations	Potential contribution
factory QA	at Bechtel, Fluor, Worley, Technip Energies, McDermott, Saipem, pressure-vessel yards, shipyards, and advanced welding/NDE suppliers.	production flow, quality, pressure boundaries, digital traceability, factory acceptance, and site installation.
Standards, licensing, and safety	ASME, ANS, ASTM, ISA, IEC, NFPA, API, IAEA, nuclear regulators, chemical process safety specialists, classification societies, and emergency-response organizations.	Open interface standard, code cases, graded QA, functional safety, hazardous areas, cyber, relief, fire protection, and responsibility boundaries.
Cost, lifecycle, and markets	NREL, NETL, Argonne, INL, DOE H2A teams, independent cost engineers, fertilizer-market specialists, water and lifecycle analysts, and university energy-system groups.	Reproducible economics, fossil benchmarks, product markets, financing, capacity factor, oxygen credit, emissions, water, and uncertainty.

Selection principles

- Select panels and contractors for evidence and capability, not prestige or advocacy for nuclear, hydrogen, ammonia, or a proprietary process.
- Include operators and maintainers who have run electrolyzers, ammonia plants, heavy transport, ports, ship loading, chemical storage, and nuclear plants—not only designers.
- Require at least one skeptic or alternative-technology expert in each major review.
- Verify access to needed facilities, data, software, process licenses, class approvals, insurance, and qualified personnel.
- Avoid panels dominated by vendors whose revenue depends on advancing to the next phase.

Scientific and Engineering Advisory Panel

Prompt to AI: If a scientific and engineering panel were established, who might serve, what expertise should be represented, and how should it operate?

Response From AI:

The panel should be small enough to deliberate, broad enough to challenge the reference concept, and independent enough to report uncomfortable findings. Names are illustrative; availability, employer approval, nationality/export restrictions, and conflicts require confirmation. Comparable substitutes are acceptable.

Expertise seat	Illustrative profile	Primary review questions
HTR-PM / HTGR reactor expert	Senior Tsinghua INET, JAEA, IAEA, or independent HTGR engineer	Reactor utility envelope, safety features, transients, data access, and limitations of the reference concept.
Nuclear integrated energy systems	Senior INL or comparable national-laboratory systems engineer	Heat/electricity coupling, dynamic operation, intermediate loop, storage, and integrated test strategy.
Nuclear licensing and QA	Former regulator or independent licensing lead with advanced-reactor and supplier-QA experience	Licensing boundary, codes, graded QA, controlled information, and regulator engagement.
Chemical process safety	Independent CCPS/HAZOP/relief/fire-protection expert	Hydrogen, oxygen, ammonia, high pressure, hazardous areas, siting, containment, and emergency response.
Hydrogen / HTSE	Independent electrolyzer scientist or operator with long-duration solid-oxide experience	Stack degradation, steam quality, system efficiency, dynamic limits, manufacturing, and pilot acceptance.
Ammonia process	Independent ammonia licensor/operator or university expert	Air separation, synthesis loop, catalyst, compression, heat recovery, refrigeration, storage, and emissions.
Mechanical and heat-exchanger codes	ASME pressure-boundary and compact-heat-exchanger expert	Materials, welds/diffusion bonding, inspection, creep-fatigue, leakage, and code qualification.
Heavy transport structural dynamics	Independent heavy-lift engineer or academic	Platform/truss loads, active jacking, carrier synchronization, route geometry, braking, and recovery.
Naval architecture / marine warranty	Independent class or marine warranty professional	Vessel stability, deck strength, ballast, lashing, weather limits, dock transfer, and voyage risk.
Modular factory construction	Independent module-yard or EPC manufacturing leader	DFMA, production takt, pressure testing, QA records, transport restraints, and site installation.
Techno-economics /	Independent NREL/NETL/university	Comparable boundaries, model replication,

Expertise seat	Illustrative profile	Primary review questions
lifecycle	cost and lifecycle expert	uncertainty, emissions, water, markets, and sensitivity.
Human factors and operations	Plant operator, maintenance, controls, and human-factors expert	Operability, alarms, procedures, staffing, maintenance access, training, and recovery from faults.

Panel charter

- Review mission, alternatives, requirements, architecture, interface standards, hazards, prototype designs, test evidence, cost models, and gate readiness.
- Receive direct access to contractors, facilities, raw data, change logs, nonconformances, model code, and owner staff.
- Issue written findings graded by safety, technical, economic, schedule, data, and transition significance; management must respond and track closure.
- Meet quarterly plus major design, hazards, test-readiness, failure, and Phase II reviews.
- Publish reports and dissenting views except for narrowly justified restricted information.
- Do not direct daily engineering, negotiate contracts, or substitute for the accountable Program Manager and safety authorities.

Management Advisory Panel

Prompt to AI: What management-oversight expertise is uniquely needed for this nuclear-integrated industrial-processing program beyond the common Management Advisory Panel charter?

Response From AI:

The Moonshot Common Framework governs the panel charter, independence, audit access, conflicts, escalation rights, reserve review, corrective action, and Managing Organization oversight. This initiative adds several seats or equivalent capabilities because it combines nuclear/process interfaces, hazardous chemical systems, marine ammonia, giant-module logistics, open standards, and international industrial supply chains.

Initiative-specific seat	Contribution	Why this program needs it
Retired nuclear-industrial cogeneration / megaproject executive	Integrated nuclear-industrial delivery	Challenges the MO on reactor-to-process boundaries, owner capability, construction/commissioning interfaces, shared infrastructure, and whether Phase II plans are executable rather than only technically interesting.
Independent process-safety / chemical-operations leader	Hydrogen, ammonia and hazardous process governance	Reviews PHA/HAZOP quality, relief/containment, hazardous inventories, operating discipline, emergency response, incident learning, and the independence of stop-work authority.
Marine ammonia / classification / port executive	Ship and waterfront deployment	Reviews ammonia storage/bunkering, marine fuel-cell demonstrations, shipboard integration, class/warranty interfaces, port handling, and whether land-based test evidence translates to marine service.
Heavy-module logistics / rail / tank-construction executive	Large infrastructure and mechanized construction	Reviews 12×96 m transport, ship/shore transfer, electrified construction rail, automated thermal-storage tank construction, maintainability, logistics bottlenecks, and shared-asset reuse.
Industrial standards / interoperability leader	Open interfaces and multi-vendor deployment	Reviews whether mechanical, thermal, piping, electrical, controls/data, cyber, rail, storage, and waterfront standards are implementable by multiple independent suppliers and supported by credible conformance tests.
China / international industrial procurement and market expert	Supply chain, integrity and transition	Challenges supplier concentration, related parties, export/sanctions constraints, localization, process-licensor dependencies, second sourcing, offtake assumptions, and successor access to sponsor-funded work.

All other management-panel seats and responsibilities are supplied by the common framework. Technology-specific scientific and engineering review remains with the Scientific and Engineering Advisory Panel in this proposal.

Agreements

Prompt to AI:

Which agreement provisions should remain inside this proposal after the portfolio-wide Funding Source-to-Managing Organization and Managing Organization-to-Recipient master terms are moved to the common framework?

Response From AI:

The generic master agreement terms are specified in the Moonshot Common Framework are not repeated here. Final contracts would use those common terms plus the program-specific schedules below, adapted by counsel to the parties, jurisdiction, nuclear/process/marine requirements, and the actual procurement instrument.

Program-Specific Schedules to the Common Agreements

Schedule	Initiative-specific contents
A - Program / Statement of Work	31 controlled 12 m × 96 m platform families; eleven baseline work packages; dual PWR/HTR-PM planning assumptions; hydrogen/ammonia and marine fuel-cell R&D; shared logistics/rail/tank/standards systems; non-nuclear prototype campus; exclusions and nuclear-interface boundaries.
B - Budget / Payment	$800M baseline technical program plus $200M centrally controlled reserve; annual profile; platform-family and shared-infrastructure cost tags; explicit shared-asset accounting with the nuclear-construction initiative; initiative-specific acceptance holdbacks.
C - Milestones / Stage Gates	Requirements/feasibility through Month 12; preliminary design and risk reduction through Month 24; detailed design/procurement through Month 36; prototype qualification through roughly Month 46; integrated/endurance demonstrations through roughly Month 57; Month 60 Phase II decisions, with package-specific safety hold points.
D - Data / IP / Open Release	Broad reusable rights in sponsor-funded platform CAD, interfaces, software, digital schemas, conformance tools, test data, materials/corrosion records, manufacturing/maintenance packages, and economics; explicit schedules for pre-existing process-licensor, reactor-vendor, or other lawful background restrictions.
E - Safety / Quality / Environment / Security	Nuclear/industrial isolation; chemical/process safety; hydrogen/ammonia hazardous areas; pressure/high-temperature systems; marine and waterfront safety; electrified rail; heavy transport; tank welding/NDE; environmental permits; cybersecurity; incident reporting; stop/hold authority.
F - Testing / Validation / Acceptance	Efficiency, heat and mass balances, temperature/pressure, throughput, emissions, corrosion/materials, catalyst/stack degradation, availability/reliability, maintenance, interoperability, transport and ship/shore handling, rail high-power behavior, tank build quality, cost/emissions reconstruction, and independent witnesses/reviewers.

G - Transition / Closeout	Per-platform Phase II packages; standards v2.0 and conformance assets; final source files/data; tooling, spares and maintenance records; supplier continuity/second source; shared-asset disposition; controlled-data disposition; first-site/first-ship recommendations and explicit stop/redirect decisions.

No duplicate master terms

Procurement, milestone-payment logic, audit, conflicts, open-source/data rights, background-IP disclosure, reserve, shared assets, changes/claims, remedies, termination, step-in, succession, insurance, and general reporting provisions are governed by the common framework unless a written program-specific schedule expressly adds or overrides a term for this initiative.

Key Codes, Standards, and Data Sources

Prompt to AI: Provide a concise list of user-provided source material, codes, standards, and cost/lifecycle tools not already summarized in the Top References chapter.

Response From AI:

Primary user-provided source material

- Climate Moonshot Business Plan, especially “Nuclear Power Climate Strategy,” “Automate Nuclear Power Construction,” and “Co-locate Chemical Processing with Nuclear Power” (google drive).
- Proposed $1B Nuclear Power Automated Site-Construction Initiative (google drive)
- Formatting instructions and gsw_research_fund_template_v4 — User-provided instruction and template files (google drive).

Nuclear process heat and technical data

- IAEA: Hydrogen Production Using Nuclear Energy.
- IAEA: Opportunities for Cogeneration with Nuclear Energy.
- INL/OSTI: High-temperature steam electrolysis reports, integrated energy systems reviews, and nuclear hydrogen cost studies.
- Generation IV and HTGR knowledge bases on process heat, intermediate loops, compact heat exchangers, materials, and safety.

Transport, marine, construction, and standards

- DNV-ST-N001 Marine Operations and Marine Warranty; applicable class rules from DNV, ABS, Lloyd’s Register, or Bureau Veritas.
- IMO Code of Safe Practice for Cargo Stowage and Securing and Cargo Securing Manual guidance.
- ASME Boiler and Pressure Vessel Code, B31 piping codes, compact heat-exchanger code-development work, and relevant ANS/ASTM/ISA/IEC/NFPA/API standards.
- Construction Industry Institute and major EPC modularization guidance; heavy-lift carrier, SPMT, rail, bridge, port, and route-engineering references.

Cost and lifecycle tools

- DOE H2A Production Models and current electrolyzer-system target definitions.
- NREL H2NEW research, manufacturing, durability, and cost publications.
- NETL, NREL, INL, and Argonne techno-economic and lifecycle frameworks for hydrogen, ammonia, electricity, heat, water, and industrial products.