




OPEN

Compute Project

White Paper: Reference Designs for Data Center Heat Reuse for Fresh Water Production and Facility Water Conservation

Revision 0.6

Last updated 2026-September 29

Revision History

Revision	Version	Date	Comments
0.0	0.0	2026/January	Executive summary & block diagrams
0.0	0.3	2026/May	Business considerations
0.0	0.4	2026/June	MEE components; water metrics
0.0	0.5	2026/July	Revise options for MEE, including heat for mechanical chiller condensers
0.0	0.6	2026/September	Business case for MEE (competitive with RO)

Authors

Authors	Company	Content type
Steven Schon	QuantaCool, USA	Original; writing; diagrams; simulations; MEE component details
Gabriel Lazar Garrido	Submer, Spain	Research; coordination; writing; water metrics

Contributors

Authors	Company	Content type
Eric Yang	My Energy LLC	
David Gyulnazaryan	Impleon	Reusing heat from mechanical chiller (condensers) to drive MEE

Contributors License



This work is licensed under a [Creative Commons Attribution-ShareAlike 4.0 International License](https://creativecommons.org/licenses/by-sa/4.0/).

Copyright info

Copyright Disclaimer under Section 107 of the Copyright Act 1976, allowance is made for "fair use" for purposes such as criticism, comment, news reporting, teaching, scholarship, and research. Fair use is a use permitted by copyright statute that might otherwise be infringing. Non-profit, education or personal use tips the balance in favor of fair use.

Table of Contents

1. Executive Summary
 2. Compliance with OCP Tenet's
 3. Introduction
 4. Business Considerations
 - a. Stakeholders
 - b. Data center heat reuse maturity model
 - c. Siting criteria for desalination
 - d. Sensitivity to data center cooling requirements
 - e. Comparison to other desalination techniques & market size
 5. Design Considerations
 - a. General design of the heat recovery station - Multi Effect Evaporation
 - b. Components and connections of the heat recovery station
 - c. Fresh water generation
 - d. Facility water conservation
 - e. Location of the heat reuse station
 - f. Power, network, cooling connections, security level for the heat reuse station
 - g. Expected redundancy/resiliency of heat reuse station
 - h. Loops design, cascading
 - i. Data center retrofits for heat reuse and water conservation/production
 - j. Environmental Impact
 6. Matching of the cooling demand
 - a. Seasonality and intra-day variations
 - b. Integration with auxiliary cooling systems
 7. Metrics
 - a. Fresh Water Generation
 - b. Energy Savings
 - c. Site Water Savings
 - d. Reduced Carbon Footprint
 8. Conclusion
 - a. Technical Feasibility and Maturity
 - b. Next steps: Integrating into Economics
 9. Abbreviations
- Appendix 1** – Reference designs, block diagrams for direct waste heat reuse from [TCS](#) to desalinate seawater.
- Appendix 2** – Reference designs, block diagrams for indirect waste heat reuse from mechanical chillers to desalinate seawater.


About Open Compute Project

1. Executive summary

This paper provides a roadmap for organizations looking to implement data center heat reuse to produce fresh water or conserve facility water. It provides technical details as well as some business and stakeholder considerations. This approach is potentially compatible with data centers' round-the-clock thermal profile.

The technology considered in this paper is multi-effect evaporation (MEE), whereby the rejected heat is displaced from conventional cooling and is used to make fresh water from impure sources (e.g., desalination) while providing the requisite cooling of the data center technical cooling system (TCS). MEE allows the heat to be reused multiple times to increase water make (Gain Ratio). MEE can yield $\sim 33 \text{ m}^3/\text{day}$ per stage per MW heat; the number of stages depends on the temperature difference between the TCS return and the impure water source. A schematic of the concept is shown below.

Heat Reuse to Produce Fresh Water via Multiple-Effect Evaporation (MEE)



Reference designs are provided for three MEE configurations to produce fresh water from coastal seawater or inland cool impure water sources such as river water, saline deepwells, or cooling tower blowdown. These designs use the waste heat from the data TCS directly, or indirectly from mechanical chillers used to cool the TCS:

- TCS heat reuse instead of rejection to dry cooler.
- TCS heat reuse instead of rejection to facility water.

- Heat reuse from mechanical chiller (MEE instead of condenser).

The objective of the reference designs is to define common features of the basic equipment and guide engineers, data center operators, facility designers, and decision-makers when designing and developing a data center to facilitate the reuse of rejected heat to make fresh water and/or minimize facility water use.

Business considerations include energy and local water availability, their value, and integration with data center operations for both new facilities and retrofits. Key stakeholders are identified, emphasizing the importance of collaboration in driving these initiatives forward. [Based on vendor presentations, MEE has competitive capital costs per unit fresh water capacity, and, assuming the waste heat is “free”, lower operating costs, than desalination via reverse osmosis \(RO\). This indicates that in water-stressed regions, desalination via data center waste heat reuse using MEE is economically competitive vs. stand-alone RO installations.](#)

The paper addresses the challenges of matching fresh water production and facility water conservation to data center cooling demand, including seasonality and intraday variations, with technical features that should be considered to cope with these variations.

The final sections identify key performance metrics and how heat reuse systems could impact them.

2. Compliance with OCP Tenets

Openness

This document is intended for widespread public use by any person interested in data centers and the reuse of their heat. Its target audience is data center developers and owners looking to evaluate the economics behind data center heat reuse. It was designed to accommodate the needs of a wide range of users, including those who need quick answers and do not normally perform “coding” or the like, and sets a framework for Developers who can work with the Python programming to contribute even deeper analysis and more features. All of the data and programming files are fully transparent.

Efficiency

The objective of this white paper is to explain the intention, the elements, and the usage of the Heat Reuse Economics tool, which has been developed to help data center owners to better understand how to reach heat reuse integration potential. Reuse of data center heat not only makes data centers more efficient but also makes off-takers who use the heat cleaner and more efficient.

Impact

This white paper and its associated tool should be used as a way of self-assessing the economics (the TCO and potential gains) of a data center heat reuse project. Optimally this tool will be used by as many data centers as possible to show the public and the authorities that they have taken measures to assess the viability of data center heat reuse from an early stage but also in the case of existing data centers.

Scale

Project success experiences, especially in Europe, demonstrate data center heat reuse is viable. This white paper and its associated tool show how a single data center can integrate any small excess heat into a heat reuse project but also applies to large data centers.

Sustainability

Data center heat reuse improves the sustainability of both data centers and end users of data center heat. For data centers, reusing their heat means reducing the amount of electricity needed for cooling and the Scope 2 emissions that may be associated with it. For heat off-takers, data center heat reduces their greenhouse gas emissions because they can replace the use of fossil fuels to produce heat with much cleaner heat. Data center heat reuse may also achieve other sustainability goals such as reduced water use.

3. Introduction

Almost all of the energy used by servers in a data center is turned into heat. Heat reuse is defined as the beneficial use of any recovered heat from the data center. To make a data center heat reuse project possible we need a combination of the following (credit: Data Centers Heat Reuse 101):

- A heat host (also referred to as the heat end user, heat off-taker, heat receiver, or heat consumer) nearby the data center area, possibly adjacent to it. If this is not the case, the collaboration of a utility will be needed as medium (or as heat host itself).
- A temperature level that is interesting for the heat host: optimally the heat will be directly used without transformations.
- Ideally the heat host and the data center have the same ownership structure; if not, then different entities must be willing to collaborate.
- Finally, a legislative structure that encourages collaboration.

This paper is part of a series of white papers that the Open Compute Project Foundation Heat Reuse group is publishing. The first white paper, [*Data Centers Heat Reuse 101*](#) was published June 2023 and describes why now is the time for heat reuse and what is needed for wide-scale adoption to turn heat from a liability into an asset that potentially can be a source of income and assist in meeting sustainability goals. A second white paper, [*Reference Designs for Data Center Heat Reuse*](#), was published March 25, 2025 and focused on reference designs defining the basic technical room in a data center for heat reuse purposes. Material has been elaborated based on community-driven input during our monthly Heat Reuse Reference Design Workstream calls.

“Basic” heat reuse systems use the relatively low-temperature (< 40 °C) rejected heat from the data centers as-is, i.e., transferring heat only to off-takers can use low-grade / mild temperature heat. Examples include on-site building heating, greenhouses, swimming pools, and preheating water for commercial laundries. Unfortunately, this severely limits the number and type of potential off-takers, and, often the heat is not needed in warmer weather, reducing the payback on the heat reuse investment.

Alternatively, a much larger pool of potential heat off-takers can be accessed if the quality (temperature) of the rejected heat can be boosted to 70+ °C, e.g., via heat pumps or supplemental heat input. Examples include district heating, food/beverage processing/ disinfection, and ab/adsorption refrigeration systems. Even so, many of these types of heat off-takers still have only intermittent or seasonal demand.

Further, a growing concern for both data centers and their surrounding communities is water consumption. Many data centers consume water for their cooling systems, either

directly (for evaporative cooling to minimize mechanical chilling and its associated power use), or indirectly (e.g. via evaporative cooling towers), to limit facility water temperatures.

To maximize economic return on heat reuse, it is preferable to have heat off-takers that can use the (un-boosted) heat rejected from the data centers continuously, round-the-clock and all year round. And, the heat off-takers should be able to tolerate short-term variability in the available heat, as data centers often see swings in server utilization. Heat reuse configurations that allow data centers to minimize their water consumption provide additional benefits to many stakeholders and broaden the acceptance of data centers by the adjacent communities.

One class of heat reuse technologies that allows year-round off-taking of heat and can tolerate variability in heat supply is evaporation (under partial vacuum) to produce fresh water from cool impure sources or brines, and displace the water consumed by conventional cooling systems. Examples of suitable water sources include seawater, brackish or river water, and cooling tower blowdown.

An option that was considered but rejected is to use heat pumps to drive desalination while shifting heat rejection from cooling towers to waterless (dry) cooling - calculations showed that the energy required by the heat pumps significantly exceeded the power required for conventional reverse-osmosis (RO) desalination.

Where data centers are in proximity to existing RO facilities, the waste heat could be used to preheat the and reduce the viscosity of the water feed to the RO systems, reducing the RO power requirements. However, the energy recovery efficiency (kW RO power reduction / kW reused waste heat) is poor (< 4%), and likely to be uneconomical.

Another option that was considered but not pursued is membrane distillation. From an energy perspective, this is the equivalent of simple single-effect evaporation, and produces less fresh water per unit of waste heat than multi-effect evaporation. Also, MD is more sensitive to fouling by contaminants in impure water feed than evaporation.

This white paper describes reference designs to use data center waste heat to recover fresh water or conserve facility water via multi-effect evaporation (MEE). It provides technical details, some business and siting considerations, and features to match fresh water production and facility water conservation to data center cooling demand. Two configurations produce fresh water by direct heat reuse from the data center technical cooling system (TCS). A third configuration to produce fresh water using the heat rejected from the data center mechanical chillers. Key performance metrics are defined and compared for examples of the various configurations.

4. Business Considerations

a. Stakeholders

- Data center operators/owners need to ensure that heat recovery and water generation don't impact operations, reliability, or security. Systems must be able to manage operations across wide and frequent variations in power loads.
- Water utilities that will manage and distribute the produced fresh water. Need to understand and quantify required additional infrastructure (e.g., piping connections, storage, disinfection) and their impact on water costs / pricing. In most parts of the world these are managed by fully public utilities (e.g. GCC, Israel, China) or PPPs (e.g. Australia, Spain). Contractors such as Veolia provide Water as a Service (WaaS) solutions that are frequently outsourced to them by public institutions.
- Water treatment facilities/maintenance waste: These vary by region but generally revolve around managing brine discharge, chemical residues, and infrastructure by-products within environmental and regulatory constraints.
- Fresh water consumers: Freshwater consumption in this context can span a spectrum of stakeholders depending on how integrated and ambitious the system design is. At the most basic level, the data center itself is the direct consumer, using freshwater for cooling, humidification, or supporting immersion systems where make-up water is required. A step further in complexity introduces industrial users within a colocation or industrial park scenario, where multiple tenants share infrastructure and water resources, which could potentially enable partial reuse loops or cascade systems (e.g., one process using the discharge from another). The most mature and complex scenario extends to supplying or integrating with the local community via municipal water networks, where freshwater flows and possibly reclaimed or thermally conditioned water are managed in conjunction with public utilities.
- Project developers & financiers, and structuring of service level agreements (SLAs). Need to address ownership and managing maintenance of the heat reuse / water purification systems (e.g., clearing clogged water intake screens/filters, cleaning of fouling deposits, maintenance of pumps and filter, etc.)

- Engineering & procurement companies (EPC) who will do detailed designs and manage installation of the heat recovery and water purification equipment. It is desirable to develop standardized and preferably modularized units.
- MEE vendors may need to modify standardized designs to operate with mild-temperature waste water sources, small differential temperatures, and corrosive effects of brines.
- Regulatory authorities. Construction, operational, and environmental permits. Incentives for implementing the heat reuse / water purification projects.
- Environmental impact assessors (government and NGOs), particularly with respect to disposition and effect of rejected (briny) concentrates on local ecosystems

b. Data center heat reuse maturity model

<similar to the corresponding section in original heat reuse reference designs white paper. Modify as needed for evaporative desalination, which is a very mature technology>

c. Siting Criteria for Desalination


To maximize return on investment, heat reuse / water purification projects are best suited for data centers in locations that:

- Are in regions where there is a shortage of fresh/potable water or the cost of water is relatively high.
- Are near cool impure water sources, e.g., coastal locations (may require water intakes below the thermocline), rivers, or cooling tower “blowdown” from power plants or heavy industry.
- Are in proximity to existing potable water infrastructure. E.g., city water mains; water treatment plants; existing conventional desalination plants
- Have sufficient open land to accommodate installation of heat reuse / water purification and ancillary equipment (e.g., heat exchangers, pumps, vessels/tankage, support structures and piping).
- The technical cooling systems (TCS) of the data center uses liquid- or 2-phase cooling. The cooled TCS coolant temperature returning to the IT side is at least 10°C warmer than the impure water source supply temperature. (Allows an additional evaporator stage or effect for every 5°C

temperature difference). Alternatively, if the data center uses a mechanical chiller (e.g., to cool facility water or data room air), the chiller's refrigerant condenser exit temperature is at least 10°C warmer than the impure water source supply temperature.

d. Comparison to other desalination techniques

Metric	Reverse Osmosis	Conventional MEE	DC-MEE
<i>Net energy cost</i>	Low	Medium	– (recovered)
<i>Primary energy</i>	++ (pumping)	++++ (thermal & pumping)	+ (pumping)
<i>Water Output (m3/day)</i>	5000-6500	400-600	30-160 per MW waste heat (1-5 effects)
<i>Energy to Output (kWh/m3)</i>	2-5	8-15 ¹	“Free” heat



Reverse Osmosis: Energy consumption per freshwater produced²

e. Economic justification & market size

The potential markets for data center waste heat desalination are coastal regions that are water-stressed. An AI query found there is over 3400 MW of known

¹ <https://www.sciencedirect.com/science/article/pii/S2590123026003828#bib0045>


² [Cost Studies of Reverse Osmosis Desalination Plants in the Range of 23,000–33,000 m3/day](#)

existing data center capacity > 10 MW within ~20 km of an ocean/sea, in water-stressed areas, and over 9800 MW of planned/proposed new capacity in these regions³ - see summary and maps, below. (Note that these figures understate the true market, as no data center capacity data were found for about half of the regions).


Coastal Data Centers x Water Stress — MW Capacity Summary					
Regions with >10MW data center capacity within ~20km of an ocean/sea, in water-stressed areas. Compiled from public reporting, Aug 2026.					
Region	Country	Existing MW	Planned/Proposed MW	Total MW (Existing + Planned)	Data Confidence
Singapore	Singapore	780	300	1,080	Existing: Confirmed; Planned: Confirmed
Dubai / Abu Dhabi	UAE	n/a	n/a	0	Existing: Estimated (derived, low confidence); Planned: Not cleanly quantified
Manama	Bahrain	n/a	n/a	0	No confirmed MW data
Doha	Qatar	230	n/a	230	Existing: Confirmed (partial)
Jeddah / NEOM	Saudi Arabia	n/a	1,690	1,690	Existing: Conflicting sources; Planned: Confirmed (composite)
Tel Aviv / Haifa	Israel	n/a	n/a	0	No confirmed MW data
Chennai	India	191	432	623	Existing: Confirmed; Planned: Confirmed
Mumbai	India	595	1,518	2,113	Existing: Confirmed; Planned: Confirmed
Sydney	Australia	770	1,889	2,659	Existing: Confirmed; Planned: Confirmed (mixed certainty)
Perth	Australia	n/a	200	200	Planned: Confirmed
Cape Town	South Africa	111	580	691	Existing: Confirmed; Planned: Confirmed
Sines (Lisbon region)	Portugal	14	1,200	1,214	Existing: Confirmed; Planned: Confirmed
Barcelona	Spain	n/a	n/a	0	Planned: Not city-specific
Hong Kong	Hong Kong SAR	n/a	920	920	Planned: Confirmed
Corpus Christi	USA (Texas)	0	704	704	Existing: Confirmed (zero); Planned: Confirmed
Tianjin / Qingdao	China	n/a	n/a	0	No confirmed MW data
Batam	Indonesia	n/a	450	450	Planned: Confirmed (partial)
Adelaide	Australia	n/a	n/a	0	No confirmed MW data
Casablanca	Morocco	n/a	n/a	0	No confirmed MW data
Montevideo	Uruguay	n/a	n/a	0	No confirmed MW data
Dublin	Ireland	738	n/a	738	Existing: Confirmed; Planned: Not cleanly quantified
Alexandria (coastal governorate)	Egypt	n/a	n/a	0	No confirmed MW data
Oahu (Honolulu)	USA (Hawaii)	6	n/a	6	Existing: Confirmed (below original >10MW screening bar)
San Juan	Puerto Rico (USA)	2	n/a	2	Existing: Confirmed (below original >10MW screening bar)
TOTAL		3,437	9,883	13,320	
Note: 'n/a' = no reliable public MW figure found; treated as 0 in totals via N(). Totals therefore understate the true market — see 'Notes & Methodology' tab. UAE (~360MW derived estimate) and Saudi Jeddah/NEOM existing (conflicting sources) are excluded from Existing MW; UAE planned (950MW total-by-2028 target) and Barcelona/Spain planned (600MW Iberia-wide, not city-specific) are excluded from Planned MW. Oahu (5.7MW) and San Juan (2MW) are included in this tally at user request even though they fall below the >10MW screening bar used for every other row — see their Confidence notes.					

Existing Data Center Capacity > 10 MW in water-stressed coastal regions

³ Claude.AI session ID cse_013xDqwpDV5EYdVK7MJYGkwo



Proposed New Data Center Capacity > 10 MW in water-stressed coastal regions



The following assumptions are conservatively made about the market for MEE desalination systems:

- As a starting point, use only the proposed new data center projects in water-stressed regions, for which capacity data is public (i.e., ignore potential for heat recovery retrofits to known existing capacity) → 9883 MW new data center power .
- Assume only 1/4 of those actually get built → 2470 MW realizable data center power.
- Assume an average PUE of 1.15 using liquid or 2-phase cooling → 2150 MW IT power

- Assume 80% of the IT power is recoverable heat (e.g. from direct-to-chip cooling) → 1720 MW recoverable heat

If the desalination modules are sized for 10 MW thermal capacity, conservatively, there is a market opportunity of over 170 MEE modules.

A triple-effect MEE unit needs only ~ 240 kWh MW heat input per m³ water make → 10 MW modules can produce ~ 1000 m³/day of fresh water.

Based on a presentation by Alfa-Laval to the OCP data center heat reuse group, the estimated capex for MEE is ~1000 Euro (~ \$ 1150 USD) per m³ per day fresh water capacity⁴. This is competitive vs. reverse osmosis (RO): An RO membrane supplier estimates capex at \$ 1200 - 1800 USD per m³ per day fresh water capacity for medium large capacity (1000 - 10,000 m³/day) installations, and \$ 800 - 1200 (USD) per m³ per day fresh water capacity for large systems (10,000+ m³/day)⁵.

The Alfa-Laval presentation cited MEE electric power consumption of 1.5 - 2.9 kWh/m³ of fresh water, with minimal other operating costs. In comparison, RO power use is 2.5 - 4 kWh/m³ water, and only accounts for 40 -60 % of the RO operating costs (the balance being for periodic membrane replacement and water pre-treatment chemicals). This suggests that the MEE is at least operating cost-neutral (if not outright advantageous) vs. RO.

With competitive capex and lower opex, MEE - using “free” data center waste heat - is economically competitive against large stand-alone RO, even with annual heat recovery utilization as low as 50 - 75 % of nameplate capacity (i.e. the data center operates at partial loads), depending on the cost of electricity. An comparison of MEE vs RO fresh water costs for two regions is shown below:

Electricity @ \$ 0.08/kWh (e.g., Waco, Texas)

⁴ <https://www.youtube.com/watch?v=xkuuxri8OW4>

⁵ <https://bestamembrane.com/seawater-desalination-costs-roi-2025-insights/>

		MEE	Medium RO	Large RO
Recoverd heat input	MW	10	N/A	N/A
Capacity	m3/day	1014	2000	10,000
Unit capex	\$/m3 per day	\$ 1,140	\$ 1,800	\$ 1,200
payback time	years	3	3	3
Annualized Capex	\$/m3	\$ 1.39	\$ 1.93	\$ 1.19
Pump Power	kWh/m3	2.9	3.5	3.5
Power cost	\$/kWh	\$ 0.08	\$ 0.08	\$ 0.08
	% of total Opex	90%	40%	60%
Annual % utilization		75%	85%	92%
Opex	\$/m3	\$ 0.26	\$ 0.70	\$ 0.47
Annualized cost of water	\$/m3	\$ 1.65	\$ 2.63	\$ 1.66

Electricity @ \$ 0.28/kWh (e.g., Singapore)

		MEE	Medium RO	Large RO
Recoverd heat input	MW	10	N/A	N/A
Capacity	m3/day	1014	2000	10,000
Unit capex	\$/m3 per day	\$ 1,140	\$ 1,800	\$ 1,200
payback time	years	3	3	3
Annualized Capex	\$/m3	\$ 1.74	\$ 1.93	\$ 1.19
Pump Power	kWh/m3	2.9	3.5	3.5
Power cost	\$/kWh	\$ 0.28	\$ 0.28	\$ 0.28
	% of total Opex	90%	40%	60%
Annual % utilization		60%	85%	92%
Opex	\$/m3	\$ 0.90	\$ 2.45	\$ 1.63
Annualized cost of water	\$/m3	\$ 2.64	\$ 4.38	\$ 2.82

5. Design Considerations

a. General design of the heat recovery station - Multi Effect Evaporation


The (cooled) TCS temperature returning to the data center or a mechanical chiller's condenser outlet temperature matters as it will define the Gain Output Ratio (GOR) and therefore the amount of freshwater generated per day:

30 - 40 °C → Conservative TCS liquid cooling. Typical for chiller condensers.

45 - 50 °C → Already demonstrated in TCS immersion / warm-water loops / 2-phase direct-to-chip. At the upper end for chiller condensers

55 - 60 °C → Aggressive for TCS but R&D tests being carried out at scale.

Impractical for chiller condensers (reduces chiller efficiency)



See annex 1 for detailed reference design.

Assume an evaporator effect (stage) for every 5°C temperature difference between the cooled TCS coolant temperature returning to the IT side and the impure water source supply temperature. The maximum number of effects can be calculated using the following formula:

$$N = (T_{\text{TCS_cool}} - T_{\text{water_feed}}) / 5$$

where:

N = number of effects


$T_{\text{TCS_cool}}$ = temperature of cool TCS coolant returned to the servers, in °C

$T_{\text{water_feed}}$ = temperature of cold impure water feeding the effects (exiting the condenser), in °C

(Note: $T_{\text{water_feed}}$ typically ~ 5 °C warmer than the cold impure water source; smaller condenser temperature rises require more water flow and pumping power).


Target of GOR of 3-4, i.e. 3 - 4 evaporator effects, corresponding to a 20-25 °C source water / TCS return temperature difference.

MEE is readily simulated by process modeling software such as Aspen Plus, and can be used to calculate fresh water production and operating conditions, material and energy balances, and define specifications for equipment such as heat exchangers and pumps. An example of an Aspen model for a triple-effect MEE is shown below:



b. Components and connections of the heat recovery station


A generic MEE system flowsheet is shown in the figure below.



The following subsystems are required to connect the data center to the MEE and reuse the rejected heat to produce fresh water:

- Branches from the data center technical cooling system (TCS) connect to/from a first effect evaporator (see figure below). This includes isolation valves, and a


temperature controller to adjust the flow split between the evaporator and the basic TCS cooling system. The latter is used when the MEE is off-line or cannot deliver the full cooling load to maintain the cool TCS coolant going to the servers at the required temperature.



- A first-effect evaporator (see figure below). This consists of a vacuum-rated chamber, a heat exchanger TCS coolant on the hot side, and preheated impure water fed to and evaporating from the cold side (this may be done with spray across the exchanger, or flow through and evaporating within the exchanger, depending on the MEE vendor's particular design).
 - It is preferable to install a mist eliminator on the water vapor exit, to prevent carryover of entrained droplets of the impure water feed into the recovered fresh water.
 - A temperature controller maintains the TCS (cool) return temperature by setting the set-point of a pressure (vacuum) controller for the chamber, which acts on a control valve in the water vapor exit line to the downstream (deeper vacuum) stage. (The water vapor temperature in the chamber is the saturation temperature of pure water at the partial vacuum of the chamber).
 - The impure water feed is substantially higher than the maximum water evaporation rate, to ensure that the concentrate is sufficiently dilute to prevent saturation / precipitation of dissolved solids, and/or to limit the temperature rise of the boiling liquid. The flow may be controlled (e.g. via

a flow controller or orifice plate) at the maximum rate required at design capacity; at reduced heat input, the concentrate will be more dilute (and boil at a lower temperature) than at design heat loads.

- The heat exchanger is sized to achieve a 5°C delta-T between the (cooler) chamber vapor and the (warmer) TCS return temperature at design capacity.
- The level of concentrate should be maintained in the chamber below the exchanger via a stand pipe and seal leg, to allow any solids to settle out (and be removed during a maintenance cleaning), to prevent plugging of the drain piping.
- In the case where the TCS cooling system is 2-phase, it is preferable that TCS feed to the evaporator is vapor only (i.e., coming off a vapor/liquid separator in the TCS loop/CDU), and that the evaporator is elevated, allowing the condensed (and slightly subcooled) cool TCS cooling to return to the TCS loop/CDU via gravity drain).




- Additional evaporator effects (see figure below). Each consists of a vacuum-rated chamber, a heat exchanger with water vapor from the upstream effect condensing on the hot side, and pre-heated impure water fed to and evaporating from the cold side (this may be done with spray across the

exchanger, or flow through and evaporating within the exchanger, depending on the MEE vendor's particular design).


The number of total evaporator effects - including the first effect - depends on the overall temperature difference between the impure water supply (e.g. cold seawater) and the target cool TCS return temperature. This paper assumes one evaporator effect for each 5 °C available overall delta-T. (There may be a separate chamber for each effect (stage), or a single vessel may be used, with compartments for each effect, depending depending on the MEE vendor's particular design).

- It is preferable to install a mist eliminator on the water vapor exit of each chamber, to prevent carryover of entrained droplets of the impure water feed into the recovered fresh water.
- A differential temperature controller maintains the temperature difference between the (warmer) condensing water vapor and the (cooler) chamber temperature by setting the set-point of a pressure (vacuum) controller for the chamber, which acts on a control valve in the water vapor exit line to the downstream (deeper vacuum) stage. (The water vapor temperature in the chamber is the saturation temperature of pure water at the partial vacuum of the chamber).
- The impure water feed and control method is the same as for the first effect, as the evaporator heat duty is substantially the same as that in the first effect.
- The heat exchanger is sized to achieve a 5 °C delta-T between the (cooler) chamber vapor and the (warmer) condensing upstream water vapor temperature at design capacity.
- The level of concentrate should be maintained in the chamber below the exchanger via a stand pipe and seal leg, to allow any solids to settle out (and be removed during a maintenance cleaning), to prevent plugging of the drain piping.
- The condensed fresh water should drain from each exchanger via a seal leg, to ensure pressure separation/isolation between successive effects.



- A final-stage condenser and vacuum system (see figure below). This consists of a heat exchanger with water vapor from the final upstream effect condensing on the hot side and cold impure water fed to and being preheated in the cold side, and a vacuum pump. An excess of impure water (beyond what is fed to the effects) is fed to the condenser, to ensure that all the water vapor is fully condensed and slightly subcooled. The excess water is diverted to the concentrated impure water return system (see below).
 - A differential temperature controller maintains the temperature difference between the (warm) condensing water vapor and the (sub-cooled) exiting final fresh water condensate, acting on a flow control valve in excess water line. This minimizes the excess water fed to the condenser, maximizing the preheat of the water fed to the evaporator effects.
 - Another temperature controller maintains the temperature of the condensing water vapor by setting the set-point of a pressure controller for the vacuum pump, which acts on an air bleed to the vacuum pump suction. This ensures that the water vapor is condensable (i.e., that the vacuum is not below the vapor pressure of water at the warmest point in the condenser).


- The heat exchanger is sized to achieve 3 °C of subcooling of the final fresh water condensate at design capacity and water feed temperature.
- The condensed fresh water should drain from the condenser via a seal leg, to ensure pressure separation/isolation from upstream effects.
- The vacuum pump should be capable of pulling a deep enough vacuum to get water to boil at the cold feed supply temperature. A mechanical vacuum pump (with a redundant spare) may be used, or, alternatively, vacuum can be provided by a liquid-jet eductor whose motive fluid is impure water driven by the feed system pump.




- A feed system for the impure water (see figure below). This consists of inlet screens/filters and redundant pumps.
 - The inlet screens/filters are to exclude coarse solids (e.g., sediments, trash, sea creatures) that can damage the pumps or foul the heat exchangers. They should be redundant, to allow maintenance without shutting down the MEE system.
 - The pumps are sized to deliver the full flow to the condenser (including the excess feed (plus any additional flow to drive liquid jet vacuum educators, if used) at design capacity, with sufficient lift and discharge pressure to move the water from a submerged intake across all the downstream

equipment and controls. Pumps should be tolerant of solids that may pass the inlet screens/filters (e.g., open-impeller designs). If the pumps are sub-surface, means for access for routine maintenance or replacement should be provided.

- The pump suction is preferably positioned where the impure water is consistently cold, e.g., below the deepest seasonal thermocline in the case of seawater.
- To avoid suction lift issues, the pump should be submerged below the water surface (pressure of water column above intake offsets the pressure needed to lift water from the intake into the pump). To facilitate maintenance of the pump and filters, it may be convenient to house the pump off-shore in a tethered structure on floating pontoons, connected to the land via a pipeline (and power cables) with flex connections to accommodate variable position due to waves and tides.




- Fresh water storage and transfer pump (see figure below). This consists of a surge (storage) tank and redundant pumps.
 - The tank should be sized to hold 2 days' worth of full fresh water production, to avoid impacting data center operation in the event of an extended outage at the receiving water utility.
 - The tank is vented to the vacuum system, to prevent back flow into the evaporator condensers, which could flood the exchangers, reducing capacity or stall operation.
 - The fresh water enters the tank via a dip-pipe, to ensure a liquid seal between the tank and each effect, ensuring that all vessels maintain their respective (different) levels of vacuum. (Note: this may create small differences in liquid height in the condensate seal legs, due to the differences in vacuum between effects).
 - The pumps are sized to deliver the full fresh water production from all the evaporator effects at design capacity, with sufficient lift to move the water from vacuum to the receiving water utility facility or pipeline.



- Concentrated impure water (e.g., brine) storage and transfer pump (see figure below). This consists of a surge (storage) tank and redundant pumps.
 - The tank should be sized to hold 8-12 hours' worth of full impure water feed flow (including the excess), to avoid impacting data center operation in the event of an on-site outage (e.g. for pump maintenance).
 - The tank is vented to the vacuum system, to prevent back flow into the evaporator chambers, which could flood the exchangers, reducing capacity or stall operation.
 - The mixed concentrate and excess feed enters the tank via a dip-pipe, to ensure a liquid seal between it and each effect, ensuring that all vessels maintain their respective (different) levels of vacuum. (Note: this may

create small differences in liquid height in the concentrate seal legs, due to the differences in vacuum between effects.

- The pumps are sized to deliver the full impure water feed flow (including the excess and assuming no evaporation at design capacity, with sufficient lift to move the water from vacuum to the reject water discharge point.



- Materials of construction must take into account the corrosive effects of impure water sources such as seawater or brines. (Note that halides such as salt attack carbon steel and many grades of stainless steel).
 - Salt-resistant metals (e.g., for heat exchangers) include titanium, 5000-series aluminum alloys, and nickel-based alloys.
 - Consider the use of polymers or composites (e.g., fiber-reinforced plastic) as cost-effective alternatives to metal alloys for vessels, pumps, and piping.

c. Fresh water generation and quality

- i. Install a continuous Monitor using a GC / MAID (Micro Argon Ionization Detector) Purge and Trap technology to determine quantities of VOCs (volatile organic compounds) and other concerning substances in water.

d. Facility water conservation

Beyond producing fresh water, the MEE displaces the evaporative heat rejection path entirely. Where the TCS would otherwise reject to a facility water system cooled by wet or adiabatic towers, the first effect becomes the heat sink instead, and the evaporation

that carries the load happens under partial vacuum inside closed vessels rather than into the atmosphere. Vapour raised in each effect condenses in the next and is recovered as distillate; the only liquid stream leaving the site would be brine (see below for environmental impact). No adiabatic towers, nor primary chillers, would be used on the secondary cooling side.

The conserved volume is therefore twofold: the make-up water that would have replaced tower evaporation, plus the 5–6% blowdown associated with it. Where distillate quality permits, a share of the production can also be returned to the facility for office use, closing the site's own limited water demand. Sites in this configuration can reach a negative NWUE (see Metrics), becoming net water producers rather than consumers.

e. Location of the heat reuse station

Siting criteria (see Business Considerations) establish *where* a site qualifies. This section deals with the relative placement of the three connection points once it does: the data center, the impure water intake, and the delivery point for the product water.

The MEE station itself should be immediately adjacent to the data center. The controlling constraint is temperature: at one effect per 5 °C of available delta-T, one or two degrees lost in a long cooling system an entire effect would be lost.

The intake is the opposite case, and can sit some kilometres away. Its cost is energy, and it scales predictably. Per MW of reused heat, the final condenser rejects roughly 1 MW into the feed, requiring ~48 L/s of seawater at a 5 °C condenser rise, against ~144 m³/day of distillate at four effects an intake flow around thirty times the water produced. Intake pumping therefore dominates the MEE's power load.

For a 10 MW station (480 L/s, DN600, $v \approx 1.75$ m/s, $f = 0.015$, $\rho = 1025$ kg/m³, pump-motor efficiency 0.70), friction head is ~3.9 m per km:

Intake run	Friction head	Pump power	Energy per m ³ distillate	Share of 10 MW heat
0.5 km	2 m	14 kW	0.23 kWh/m ³	0.14%
1 km	3.9 m	27 kW	0.45 kWh/m ³	0.27%
2 km	7.8 m	55 kW	0.91 kWh/m ³	0.55%

5 km	19.5 m	135 kW	2.2 kWh/m ³	1.35%
------	--------	--------	------------------------	-------

Friction losses per Darcy-Weisbach; see Crane Technical Paper No. 410, *Flow of Fluids Through Valves, Fittings and Pipe*. Static lift, intake screen losses, and fittings are additional and site-specific. A concentrate outfall at comparable distance roughly doubles the figures unless siphon recovery is used on the return.

The practical limit is visible in the table: beyond ~2 km the intake pumping alone approaches the total energy of conventional RO, and the case for the configuration weakens. Delivery of product water is not energy-constrained, but proximity to the municipal network still matters, for pipeline capital cost, dedicated potable-standard construction, and the permitting of a new connection point.

- f. Power, network, cooling connections, security level for the heat reuse station
<Use similar wording as this section in the Reference Designs for Data Center Heat Reuse white paper ?>
- g. Expected redundancy/resiliency of heat reuse station
- h. Loops design, cascading
- i. Data center retrofits for heat reuse and water conservation/production
- j. Environmental Impact

6. Matching of the cooling demand

- a. Seasonality and intra-day variations
- b. Integration with auxiliary cooling systems

Metrics

- a. Freshwater generation

In the application presented in this whitepaper, the data center site would have a secondary benefit of behaving like a freshwater generation asset. In this case, conventional data center water metrics developed by The Green Grid and ISO/IEC 30134 become incomplete. While the Energy Reuse Factor (ERF) would quantify the reutilization of heat for any application, no metric exists for capturing the amount of freshwater coming from the facility. Water Usage Effectiveness (WUE), the main water consumption KPI during the use phase, does not capture any of the freshwater being purified from the recovered heat.

For this use case, this paper would like to introduce a derived metric called Net Water Usage Effectiveness, or **NWUE**:

$$\text{NWUE} = (W_{\text{consumed}} - W_{\text{generated}}) / E_{\text{IT}}$$

Where:

- W_{consumed} = annual on-site water consumption
- $W_{\text{generated}}$ = net freshwater exported for beneficial use
- E_{IT} = annual IT energy consumption

Reported in L/kWh_{IT}, this metric preserves the logic of WUE while allowing the result to go below zero. A negative WUE_{net} would indicate that, over the reporting period, the facility is a net water producer rather than a net water consumer.

That said, WUE_{net} should not be read on its own. It should be paired with existing heat reuse metrics (like the ERF mentioned earlier), so the industry can measure the effectiveness of the heat reuse infrastructure of a given site. For example, a strong ERF with a poor WUE_{net} may indicate that heat is being exported towards the desalination plant, but not converted into water efficiently enough to matter. This will also ensure optimal operations and return on investment for the significant additional CAPEX spent on the desalination collocation.

Another potential sub-KPI that could measure the desalination effectiveness could be Thermal-to-Water Productivity:

$$\text{TWP} = W_{\text{generated}} / E_{\text{thermal,in}}$$

Where:

- $W_{\text{generated}}$ = net freshwater exported for beneficial use, and;
- $E_{\text{thermal,in}}$ = thermal energy delivered to MEE system

Also reported in L/kWh, this would show how effectively recovered heat is turned into usable water. That is particularly relevant because the efficiency of the MED thermal process should also be monitored and improved in operations. The results of the data center-colocated MED can be compared to traditional, stand-alone MED systems. It will not necessarily require a like to like result, but a certain threshold will need to be reached to ensure viability of these projects. As this suggests, the TWP metric will be particularly important in the piloting phase.

A few boundary conditions should be explicit in the methodology. First, freshwater should be reported on a net beneficial-use basis, not gross distillate production. Any water internally consumed for polishing, cleaning, blowdown management, or auxiliary processes should be deducted before credit is claimed. Second, additional electricity loads for pumps, controls, post-treatment, and intake or discharge infrastructure should be disclosed separately, otherwise the water metric can look better than the full system really is. Third, quality class matters: high-purity distillate reused inside the facility is not the same as potable water exported for municipal or industrial use, and the reporting should distinguish between those outputs. Fourth, uptime and temporal matching matter. The value of this configuration comes from the fact that IT load, and therefore rejected heat, is available continuously, allowing MED to run as a 24/7 application rather than as an intermittent thermal sink.

If this use case proves credible and repeatable, the next logical step would be for The Green Grid to recognize and socialize the metric as an accepted derivative of existing xUE thinking, with eventual consideration for inclusion or alignment within the ISO/IEC 30134 KPI family.

b. Energy savings

- Reduced data center cooling power.

Diverting heat from the TCS primary cooling systems to drive MEE can reduce data center power consumption depending on the type of primary cooling system (see block diagrams in Appendix 1).

- Where the TCS normally rejects heat via dry coolers, the dry coolers can be turned off, saving the power used to drive the dry cooler fans.
- Where the TCS normally rejects heat via facility water service (FWS) cooled by wet cooling towers, the cooling towers can be turned down or off, saving the power used to drive the FWS pumps and cooling tower fans.

- Where the TCS normally rejects heat via FWS cooled by dry coolers, the FWS can be turned down or off, saving the power used to drive the FWS pumps and dry cooler fans.
- Where the TCS normally rejects heat via FWS cooled by mechanical chillers, the FWS can be turned down or off, saving the power used to drive the FWS pumps and chiller machinery.

The net power savings need to account for power used to operate the MEE system (feed pumps, fresh water pumps, reject brine pumps, and vacuum system).

The power savings can be used to calculate the reduction in the data center's PUE and operating costs and the value per MW of heat reuse.

<Provide example calculation>

- o Avoided stand-alone desalination power

Reusing waste heat from data centers to drive MEE avoids the power (and capital) associated with the water off-takers adding equivalent stand-alone (external) desalination capacity such as R.O. units (use electric power) or steam-driven MEE (uses fuels).

The avoided power use can be used to calculate the reduction in the water utility's operating costs and unit cost of fresh water.

<Provide example calculation>

c. Site water savings

- o Reduced data center site water consumption:

Diverting heat from the TCS primary cooling systems to drive MEE can reduce data center power water consumption depending on the type of primary cooling system (see block diagrams in Appendix 1). These should be included in the (negative) WUE calculations

- Where the TCS normally rejects heat via facility water service (FWS) cooled by wet cooling towers, the cooling towers can be turned down or off, saving the make-up water used to replace the water evaporated by the cooling towers, as well as the additional water (typically 5 - 6% of the evaporative load) used for "blowdown"

to prevent accumulation and precipitation of minerals and sediments in the FWS loop from the makeup water.

The water savings can be used to calculate the reduction in the data center's WUE and operating costs and the value per MW of heat reuse.

<Provide example calculation>

d. Reduced carbon footprint

- Reduced data center and water off-taker carbon footprint:

The energy savings by reusing the data center heat translates into Scope 2 greenhouse gas (GHG) emission reductions if the power supplied to the data center is derived from fossil fuels.

Similarly, the avoided power to drive (external) desalination systems of equivalent capacity translates into Scope 2 greenhouse gas (GHG) emission reductions if the power supplied to water utility is derived from fossil fuels.

However, the Scope 3 GHG emissions “embedded” in the materials, transportation, and construction of the MEE equipment (and avoided external desalination equipment) needs to be accounted for, to calculate the overall life cycle net GHG reduction.

Where carbon credits are available, the GHG reductions be used to calculate the reduction in the data center's operating costs and the value per MW of heat reuse.

<Provide example calculation>

7. Conclusion

8. Abbreviations


<List all the abbreviations here>

Appendix 1


Reference designs, block diagrams for direct waste heat reuse from TCS to desalinate seawater.

The waste heat can be reused directly from the TCS liquid or 2-phase coolant to drive the MEE, as shown in the block diagrams below. These are illustrative only, for the case where the TCS return temperature is 32 °C and the seawater is < 10 °C. The number of evaporative effects will depend on the available temperature difference between the TCS return and the impure water (e.g., sea water) exiting the condenser.


Waste heat reuse from liquid-cooled IT to desalinate seawater.




Waste heat reuse from 2-phase-cooled IT to desalinate seawater.




There are several ways that a data center TCS rejects heat, e.g., to dry coolers, to facility water systems that are cooled by wet or dry coolers, or to facility water systems that are cooled by mechanical chillers. The rejected heat can be diverted from these conventional primary cooling systems (in hatched boxes) to drive the MEE, as shown in the schematic diagrams below.



MEE Instead of Dry Cooling of TCS Coolant



MEE Instead of Wet or Dry Cooling of Facility Water




MEE Instead of Mechanical Cooling of Facility Water

Appendix 2


Reference designs, block diagrams for indirect waste heat reuse from mechanical chillers to desalinate seawater.

The waste heat can be reused indirectly from a TCS (air, liquid or 2-phase) that is cooled by a mechanical chiller, e.g., in regions (such as the Middle East or Southeast Asia) where ambient conditions preclude air cooling of facility water. In this case, the chiller compressor's heat that is normally rejected to the chiller refrigerant loop's (post-compressor) condenser can be used to drive the MEE, as shown in the block diagram below. These are illustrative only, for the case where the chiller condenser outlet (subcooled) temperature is 37 °C and the seawater is < 10 °C. The number of evaporative effects will depend on the available temperature difference between the chiller condenser outlet and the impure water (e.g., sea water) exiting the condenser.

Waste heat reuse from mechanical chiller to desalinate seawater.



A data center that TCS rejects heat indirectly to a chiller system typically does so via a facility water system (FWS) loop, which in turn is cooled by and transfers heat to the chiller. The chiller uses a refrigerant loop evaporating at low temperature and pressure. The refrigerant gas is compressed by an electrically driven compressor, boosting the pressure and temperature of the refrigerant. The hot high temperature refrigerant is condensed and sub-cooled in a condenser, and then subsequently flash-expanded (cooled and partially vaporized) and returned to the evaporator, closing the loop. The rejected heat can be diverted from the chiller condenser (in hatched box) to drive the MEE, as shown in the schematic diagram below.



About Open Compute Project

The Open Compute Project Foundation is a 501(c)(6) organization which was founded in 2011 by Facebook, Intel, and Rackspace. Our mission is to apply the benefits of open source to hardware and rapidly increase the pace of innovation in, near and around the data center and beyond. The Open Compute Project (OCP) is a collaborative community focused on redesigning hardware technology to efficiently support the growing demands on compute infrastructure. For more information about OCP, please visit us at <http://www.opencompute.org>