

*Series Editor*  
*Serge Petiton*


---

# **Energy-Efficient Computing and Data Centers**


---

Luigi Brochard  
Vinod Kamath  
Julita Corbalán  
Scott Holland  
Walter Mittelbach  
Michael Ott


Color section




**Figure 2.1.** *Heat flux on an air-cooled chip*



**Figure 2.2.** *Lenovo 1U full width SR630 server*



**Figure 2.3. Front view of an SR630 server**



**Figure 2.4. Rear view of an SR630 server**






Figure 2.5. Typical efficiency curve for an air moving device


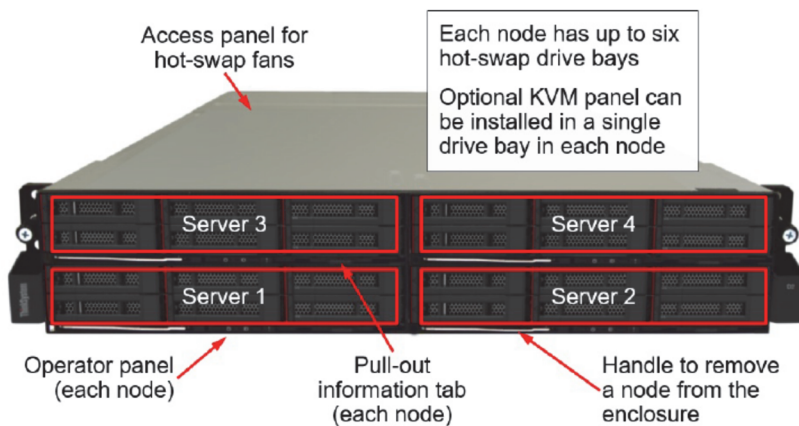




Figure 2.6. Thermal resistance and air flow pressure drop




**Figure 2.7.** Front view of the 2U Lenovo SD530 server with 4½ width nodes




**Figure 2.8.** Inside view of the ½ width node of a Lenovo SD530



**Figure 2.10.** *Thermal transfer module for SD530*



**Figure 3.1.** *Classic air-cooled machine room*



**Figure 3.2.** *Hot and cold air flows in the machine room*








Figure 3.3. ASHRAE classes of data center operation



**Figure 3.6.** *Principle of passive RDHX*



**Figure 3.7.** *RDHX, CDU and chillers*



**Figure 3.8. RDHX heat removal capacity for a 32 kW and 10 kW rack load**

## Water

1. High heat capacity

$$c_v \approx 1 \text{ Wh}/(\text{L}\cdot\text{K})$$

2. Low thermal resistance



$$\dot{q}'' = R_{th} \cdot \Delta T$$

$$R_{th} = 0.1 \text{ K cm}^2 / \text{W}$$

$$\dot{q}'' = 50\text{--}100 \text{ W/cm}^2$$

$$\Delta T = 5\text{--}10^\circ\text{C}$$

## Air

1. Low heat capacity

$$c_v \approx 0.0003 \text{ Wh}/(\text{L}\cdot\text{K})$$

2. High thermal resistance




$$\dot{q}'' = R_{th} \cdot \Delta T$$

$$R_{th} = 1 \text{ K cm}^2 / \text{W}$$

$$\dot{q}'' = 50\text{--}100 \text{ W/cm}^2$$

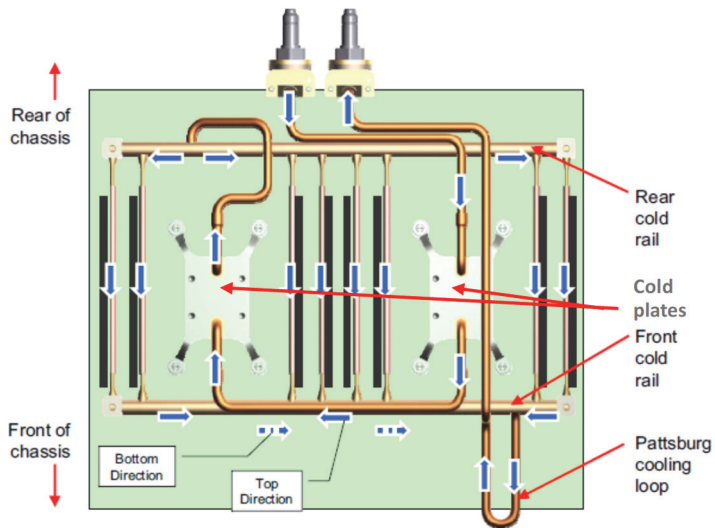
$$\Delta T = 50\text{--}100^\circ\text{C}$$

**Figure 3.9.** Water versus air heat capacity and thermal resistance




**Figure 3.11.** Water loop assembly of the IBM iDataPlex dx360 m4







**Figure 3.12.** *Water flow path in the IBM iDataPlex dx360 m4*



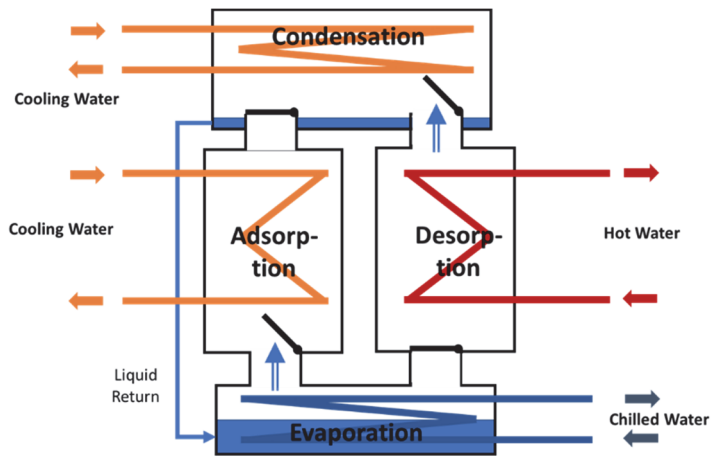
**Figure 3.13.** *Lenovo ThinkSystem SD650 water-cooled node*




**Figure 3.14.** *Lenovo ThinkSystem SD650 memory water channels*




**Figure 3.15.** Adsorption chiller principle and phases (two chambers design)




**Figure 3.16.** *Adsorption chiller principle and phases (four containers design)*



**Figure 3.17.** Sankey diagram of the energy flows in an adsorption chiller



**Figure 3.18.** Measured  $COP_{el}$  values of different cooling technologies



**Figure 3.19.** *Water uptake of new adsorbent material*




Figure 4.9. SP and DP peak GFlops per watt for NVIDIA GPUs

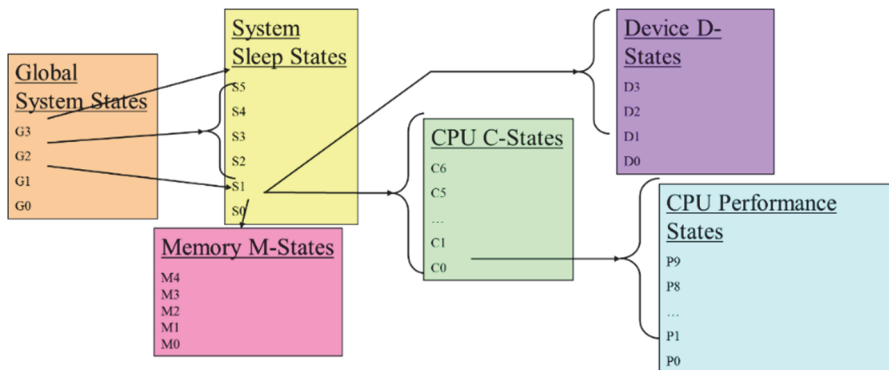


Figure 4.10. ACPI states









Figure 5.3. Node, CPU and DIMM DC power of SIMD instructions Turbo OFF




**Figure 5.5.** Node, CPU and DIMM DC power running HPL Turbo OFF




**Figure 5.6.** Node, CPU and DIMM DC power running HPL Turbo ON




**Figure 5.7.** CPU frequency and temperature running HPL with Turbo OFF




**Figure 5.8.** CPU frequency and temperature running HPL with Turbo ON




**Figure 5.9.** CPU0 and CPU1 CPI and node bandwidth running HPL with Turbo OFF




**Figure 5.10.** CPU0&1 CPI and node bandwidth running HPL with Turbo ON




**Figure 5.11.** Node, CPU and DIMM power running STREAM with Turbo OFF




**Figure 5.12.** Node, CPU and DIMM power running STREAM with Turbo ON




**Figure 5.13.** CPU temperatures and frequencies running STREAM with Turbo OFF




**Figure 5.14.** CPU temperatures and frequencies running *STREAM* with Turbo ON



**Figure 5.15.** CPU CPIs and node bandwidth running *STREAM* with Turbo OFF




**Figure 5.16.** CPU CPIs and node bandwidth running STREAM with Turbo ON




**Figure 5.17.** Node, CPU and DIMM power running BQCD1K with Turbo OFF







**Figure 5.18.** Node, CPU and DIMM power running BQCD128 with Turbo OFF




**Figure 5.19.** Node, CPU and DIMM power running BQCD1K with Turbo ON




**Figure 5.20.** Node, CPU and DIMM power running BQCD128 with Turbo ON




**Figure 5.21.** CPU temperatures and frequencies running BQCD1K with Turbo OFF




**Figure 5.22.** CPU temperatures and frequencies running BQCD128 with Turbo OFF




**Figure 5.23.** CPU temperatures and frequencies running BQCD1K with Turbo ON




**Figure 5.24.** CPU temperatures and frequencies running BQCD128 with Turbo ON




**Figure 5.25.** CPU CPIs and node bandwidth running BQCD1K with Turbo OFF



**Figure 5.26.** CPU CPIs and node bandwidth running BQCD128 with Turbo OFF




**Figure 5.27.** CPU CPIs and node bandwidth running BQCD1K with Turbo ON




**Figure 5.28.** CPU CPIs and node bandwidth running BQCD128 with Turbo ON


SD530 - Turbo Off - BQCD1K - CPU Temp & Freq




SD650 - Air-Cooled - Turbo Off - BQCD1K - CPU Freq. & Temp.



SD650 - Air-Cooled - Turbo Off - BQCD1K - CPU Temp & Freq - Swapped









**Figure 5.29.** Comparison of BQCD1K with Turbo OFF on three servers


SD530 - Turbo On - BQCD1K - CPU Temp & Freq



SD650 - Air-Cooled - Turbo On - BQCD1K - CPU Freq. & Temp.



SD650 - Air-Cooled - Turbo On - BQCD1K - Temp & Freq Swapped










Figure 5.30. Comparison of BQCD1K with Turbo ON on three servers

SD530 - Turbo On - Linpack - CPU Freq. & Temp.




— ptumon CPU0 Core Freq — ptumon CPU1 Core Freq  
— ptumon CPU0 Temp — ptumon CPU1 Temp

SD650 - Air-Cooled - Turbo On - Linpack - CPU Freq. & Temp.




— ptumon CPU0 Core Freq — ptumon CPU1 Core Freq  
— ptumon CPU0 Temp — ptumon CPU1 Temp

SD650 - Air-Cooled - Turbo On - Linpack  
CPU Freq. & Temp. Swapped




— ptumon CPU0 Core Freq — ptumon CPU1 Core Freq  
— ptumon CPU0 Temp — ptumon CPU1 Temp


SD530 - Turbo On - Linpack - Node, CPU & DIMM Power




SD650- Air-Cooled - Turbo on - Linpack - Node, CPU & DIMM Power



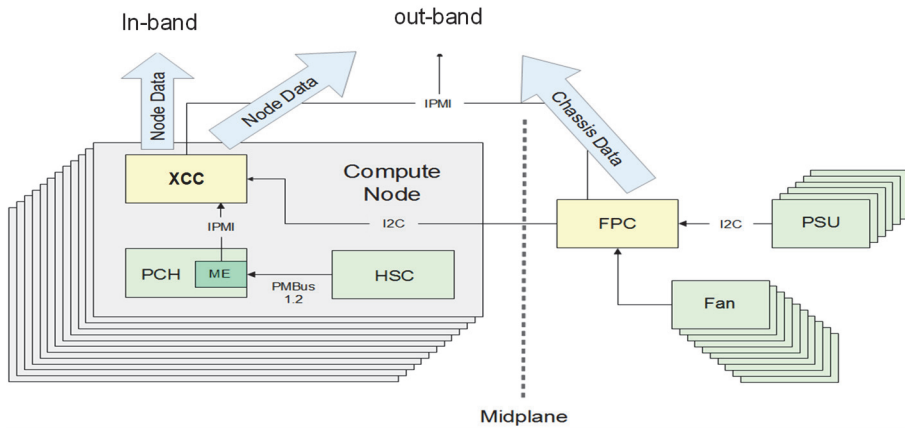
SD650 - Air-Cooled - Turbo On - Linpack - Node, CPU & DIMM Power Swapped



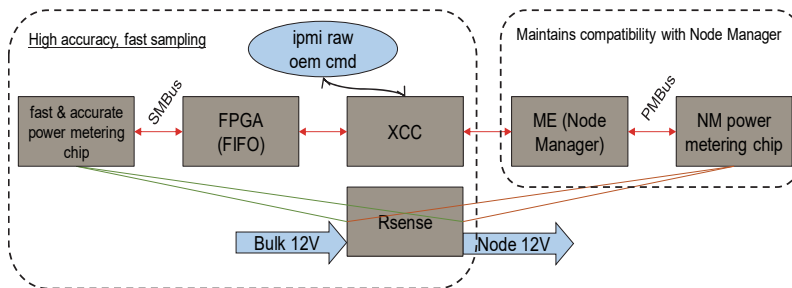
**Figure 5.31.** Comparison of HPL with Turbo ON on three servers




**Figure 5.32.** Cooling impact on 2697v3 temperature and performance



**Figure 6.1. Node power management on Lenovo ThinkSystem**




**Figure 6.4.** *New circuit to get higher resolution DC node power on SD650.*




**Figure 6.5.** *Power accuracy of different circuits*






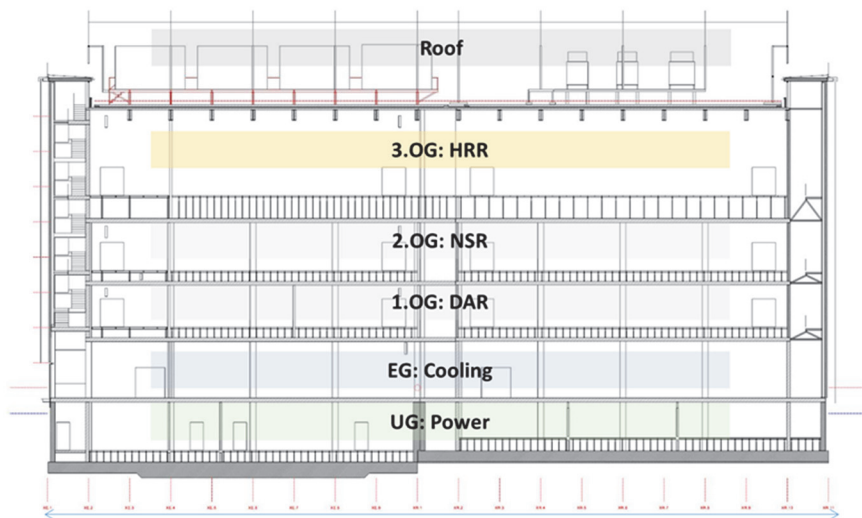
**Figure 7.1.** RSF load profile for the first 11 months of operations



**Figure 7.2.** RSF hourly PUE over the first 11 months



**Figure 7.3.** *RSF ERE as a function of outdoor temperature*



**Figure 7.4.** *Sectional view of LRZ data center*

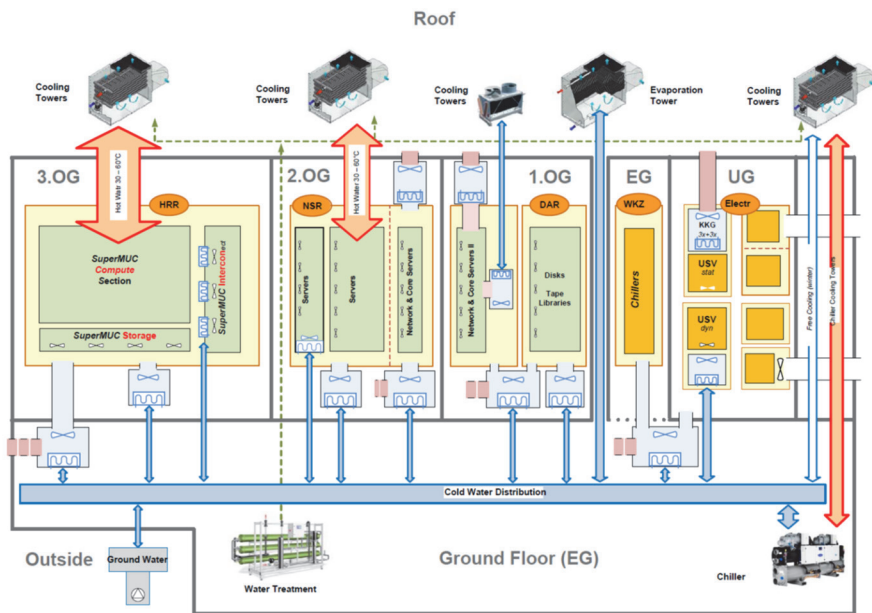



Figure 7.5. LRZ cooling infrastructure overview



**Figure 7.10.** CoolMUC-2 operations 05/2017–09/2017

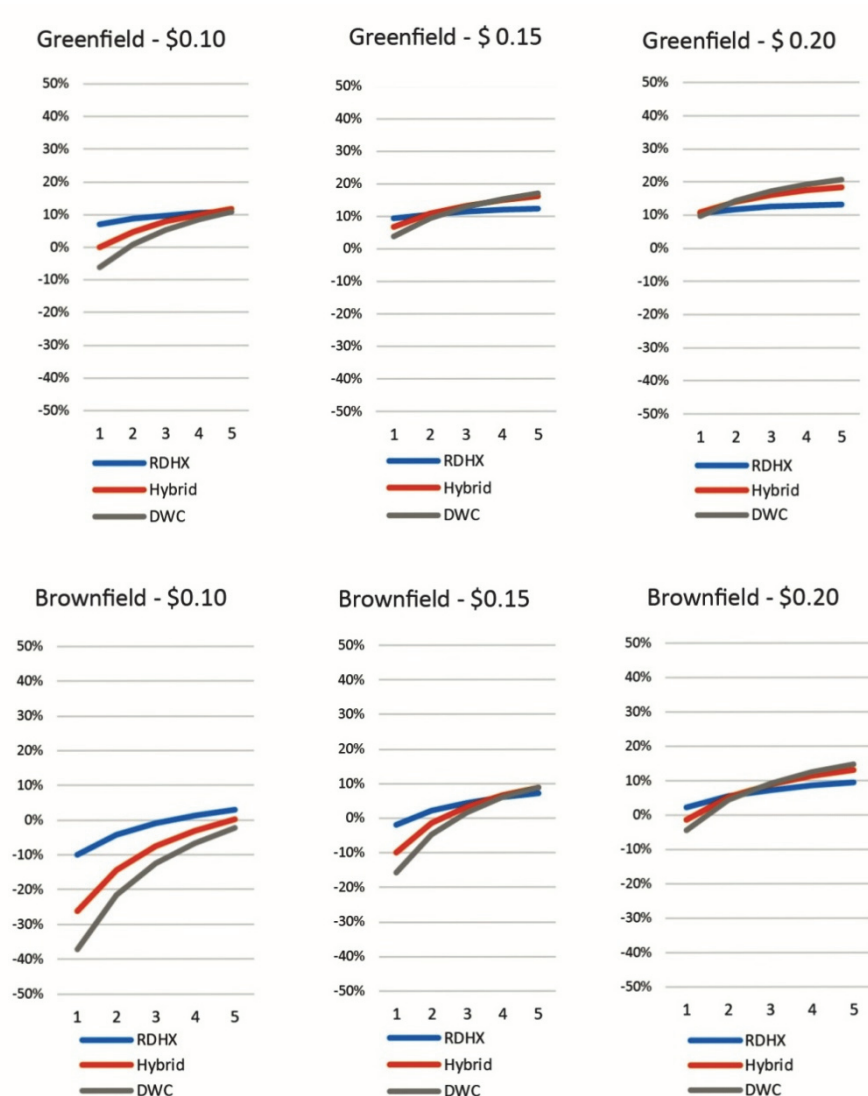




Figure 7.11. Impact of electricity on project payback


Greenfield - 150W




Greenfield - 200W




Greenfield - 250W



Brownfield - 150W



Brownfield - 200W



Brownfield - 250W




Figure 7.12. Impact of SKU TDP on project payback at \$ 0.15/kWh electricity price

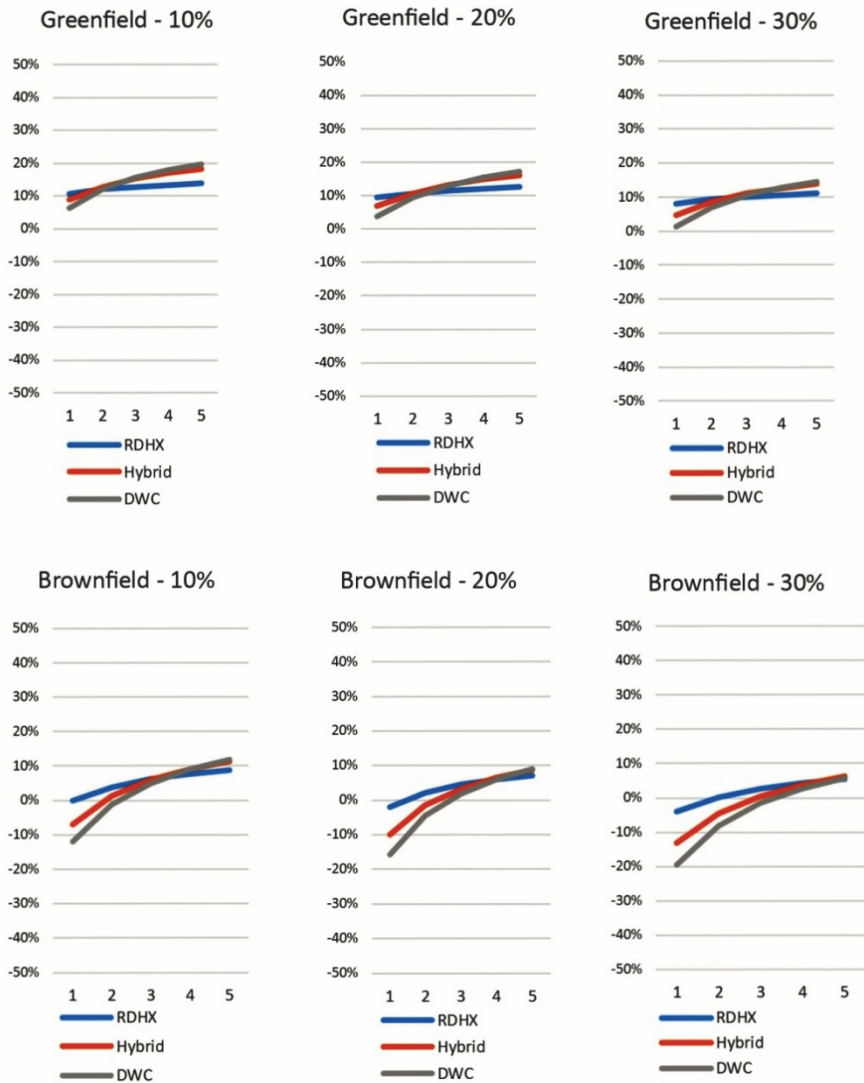


Figure 7.13. Impact of free cooling ratio on project payback at \$0.15 electricity price



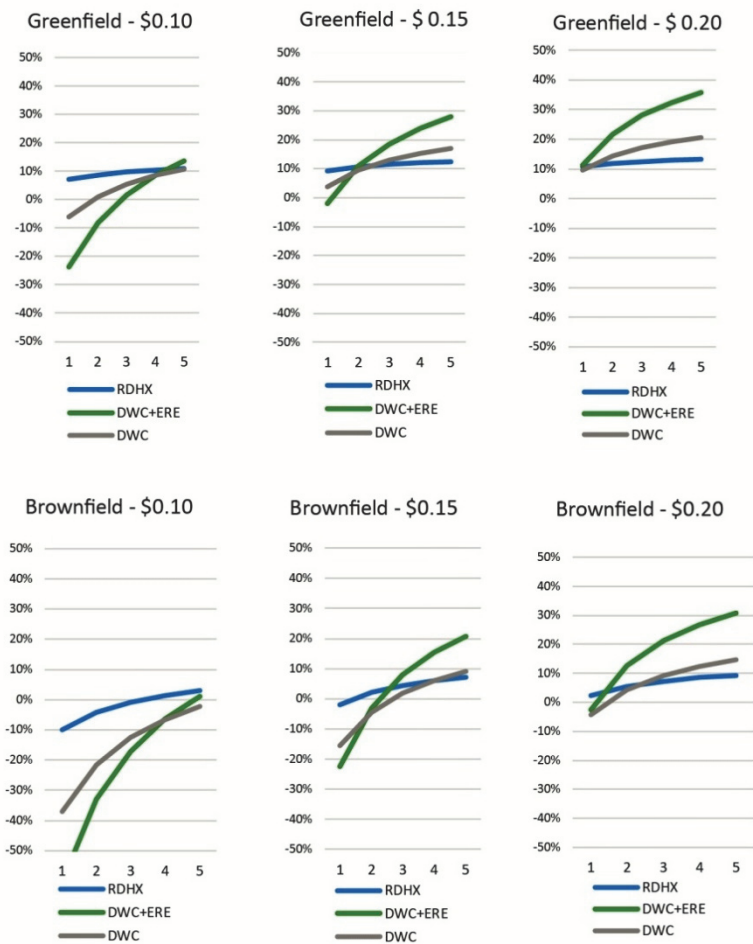
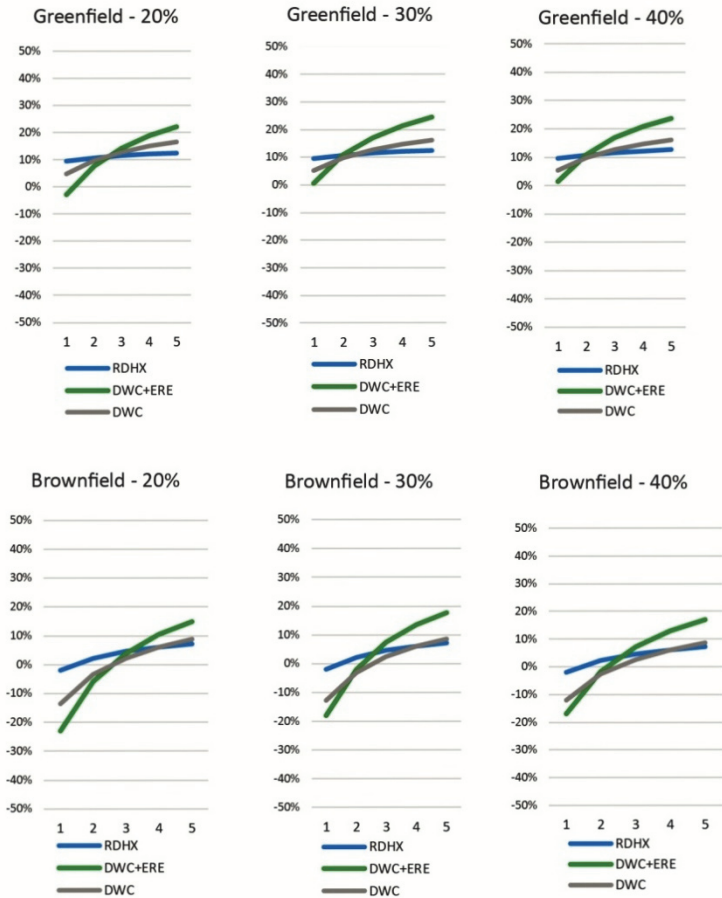


Figure 7.14. Impact of hot water reuse on project payback and different electricity price



**Figure 7.15.** Impact of hot water energy reused on project payback with various air to water power ratios, \$0.15/kWh electricity price, 20% free air-cooling ratio and 50% COP




Figure 7.16. PEMFC diagram


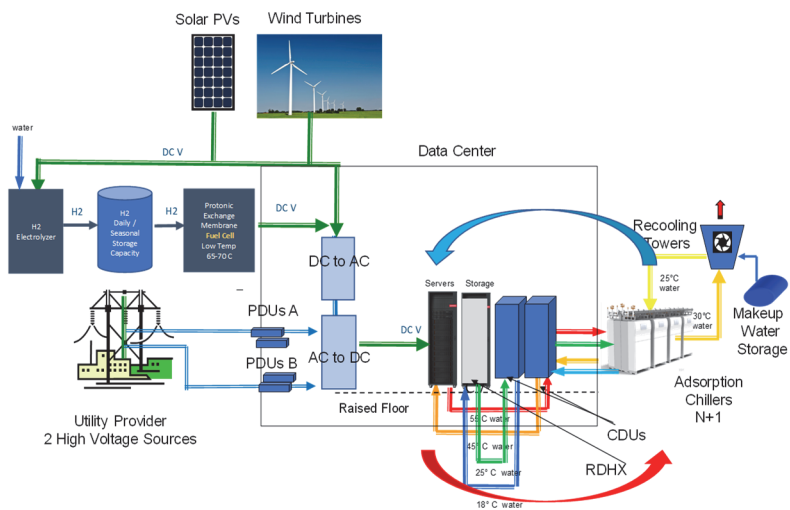


Figure 7.17. Storing and reusing excess energy with PEMFC



**Figure 7.18.** *Toward a net-zero energy data center*