

Drought and Extreme Heat Impacts to Data Centers in Northern California

Dr Gemma J. Anderson
Lawrence Livermore National Laboratory

Unclassified

July 20, 2023

Technical Report Prepared for:
Infrastructure Security Division
(ISD), Cybersecurity and
Infrastructure Security Agency
(CISA), Department of Homeland
Security (DHS)



Disclaimer

This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.

Lawrence Livermore National Laboratory is operated by Lawrence Livermore National Security, LLC, for the U.S. Department of Energy, National Nuclear Security Administration under Contract DE-AC52-07NA27344. This report is released with tracking number: LLNL-TR-852189.

Acknowledgements

The author would like to acknowledge and give a special thanks to Dr. Otto Alvarez (LLNL) who designed the geospatial system that is used to display the results presented in the case studies chapter of this report. The author is deeply grateful to Sunny Wescott (DHS) for invaluable feedback on earlier versions of this manuscript, and for providing helpful suggestions throughout the project. A sincere thank you is also extended to Andrew Bochman (INL), Christopher Reidel (DHS), Bill DeLong (DHS), and Mario Garcia (DHS), and many others that provided their expertise and guidance.

Table of Contents

Executive Summary	4
1. Introduction.....	5
2. Climate science	7
2.1 Earth’s changing climate.....	7
2.2 Climate models and emission scenarios	7
2.3 Climate impacts at a regional scale.....	8
3. Extreme heat and drought impacts to data centers	10
3.1 Assessing relevant impacts	10
3.2 Data center relevant impacts.....	12
4. Case studies.....	16
4.1 San Jose airport.....	16
4.2 Livermore airport.....	18
4.3 Napa airport	18
5. The path forward.....	20
References.....	22

Executive Summary

Commonly, data centers are designed and operated on the assumption of stationarity of the climate. Historical weather data informs the basis of their design and operations. However, as a result of climate change, this assumption is not valid and poses a great risk for data centers. In the coming decades, Northern California's drought and extreme heat risk is projected to increase significantly. This report provides a drought and extreme heat impacts assessment of data centers in this region. Expert solicitation of data center experts enabled a tailored approach to this report. We identify the climate variables and analytics relevant for assessing data center-specific impacts from extreme heat and drought and demonstrate how to access and understand future climate projections from climate models. We emphasize the importance of including future climate projections into data center design and planning. This report is an important first step towards building an effective extreme heat and drought adaptation strategy for data centers in Northern California. While our geographical focus in this report is Northern California, many findings in this report are broadly applicable.

1. Introduction

Data centers are facilities that house computing machines and related hardware components that are used by businesses to assemble, process, store and disseminate large amounts of data. In addition to being crucial assets for an organization's everyday operations, they are also vital to the functioning of modern society, supporting online services and applications we depend upon and much of our critical infrastructure, e.g., public safety, finance, healthcare, and transportation. Unexpected data center shutdowns would have widespread repercussions to society.

The computing systems and related components that data centers house generate tremendous amount of heat. The temperature and humidity levels within the data centers need to be carefully controlled to ensure that the computing systems can perform reliably and efficiently. The American Society of Heating, Refrigerating, and Air Conditioning Engineers (ASHRAE) published the Thermal Guidelines for Data Processing Environments in 2021¹. Their recommended standards for data center cooling advise servers room temperatures that ensure reliable operation and maximize efficiency and longevity of servers. ASHRAE publishes psychrometric charts that display the recommended temperature and humidity ranges for a data center. The current recommended temperature range for data center equipment is 18 to 27 degrees C (64.4 to 80.6 degrees F), and ASHRAE also provide *allowable* ranges for four different classes of data center equipment: A1, A2, A3 and A4 that vary considerably. Most data center equipment falls into class A1 or A2.

To maintain those desired conditions in a data center requires significant support infrastructure including power subsystems, uninterruptible power supplies (UPS), backup generators, ventilation and cooling equipment. Data centers are extremely energy-intensive, consuming ~ 2% of the total US electricity.² Roughly 40% of the power that data centers consume goes toward cooling them. Data centers must be cooled through water cooling, air cooling, refrigerants, or combinations of these methods. The cooling method adopted depends on the data center size and location. Water cooling can be less energy intensive than air cooling, but their water footprint is dramatically higher. For example, Google has significantly reduced its carbon footprint by using water for cooling. However, they recently divulged their water-use data and it was found to be a staggering 3.3 billion gallons in 2021 within the US alone.³

Droughts and heatwaves, like many other climate-related hazards, are becoming more frequent and more severe as a result of climate change. Many parts of Northern California faced record-breaking temperatures in 2022, and the Southwestern US megadrought that began in 2000 is the was the driest 22-year period for at 1200 years [Williams et. al., 2022]. Hot and dry conditions can have widespread impacts on built and planned data centers. For example, extreme heat can lead to increased energy use for cooling, overheating and failure of equipment, reduced efficiency, and shutdowns/outages due to heat-related power disruptions. Extreme heat can also damage the data center building infrastructure. For data centers that use water for cooling, water scarcity arising from prolonged drought could lead to decreases in cooling capacity, and operational disruptions.

¹ https://www.techstreet.com/ashrae/standards/thermal-guidelines-for-data-processing-environments-5th-ed?product_id=2212974

² <https://www.energy.gov/eere/buildings/data-centers-and-servers>

³ <https://www.watercalculator.org/news/news-briefs/google-data-center-water/>

Water-cooled data centers also are particularly vulnerable to high humidity, which can impede their ability to cool efficiently.

Data centers are typically designed and built assuming that historical weather conditions are representative of expected future conditions. Ramifications of this incorrect assumption are already apparent. In the UK, the record-breaking temperatures last summer resulted in shutdowns of Google and Oracle data centers.⁴ Last year also saw Twitter's data center taken offline by extreme heat in California.⁵ These impactful events will become increasingly common if data centers do not plan for climate change in their design and operations.

Impacts to data centers from drought and extreme heat will be non-uniform and depend on factors unique to their specific location. Northern California (specifically Silicon Valley) is home to more than 160 data centers and is the third-largest data center market in the US. Planning tools that incorporate plausible and adequate future regional climate scenarios are needed to inform infrastructure decisions and enable prioritized hardening of data center assets against increased exposure to drought and extreme heat.

This report provides an assessment of drought and extreme heat/humidity impacts to data centers in Northern California. We developed a geospatial system to enable quantitative analysis of heat and humidity to data centers in Northern California. We worked with various key stakeholders to identify and provide data center-specific variables and analytics. Since data centers often use nearby airports for their weather data, we select three airports within Northern California and present the results for these locations as case studies. Our assessment of drought impacts is mostly qualitative in nature owing to the complex non-localized nature of the impacts of drought.

The report is structured as follows. In Section 2. we introduce the changing climate, climate models, and the climate data we use in this report. In Section 3. we describe the climate variables and analytics relevant to data centers. In Section 4. we provide the case study results. In Section 5. we discuss a path forward.

⁴ <https://www.bbc.com/news/technology-62202125>

⁵ <https://www.latimes.com/california/story/2022-09-12/twitters-data-center-knocked-out-by-extreme-heat-in-california>

2. Climate science

2.1 Earth's changing climate

Throughout history the Earth's climate has always changed, but over the last 200 years, since the beginning of the industrial revolution, the changes have been dramatic. Since 1880, the Earth has warmed 1.9° Fahrenheit. Recent global temperatures have increased at a rate unprecedented in at least the last 2000 years, and the latest decade was warmer than any multi-century period for 125,000 years [Arias et. al., 2021]. The primary driver of climate change is human activity, mainly through burning of fossil fuels such as oil, coal, and gas. Fossil fuels release greenhouse gases (GHGs) into the atmosphere that cause the Earth to warm. The effects of climate change extend beyond increasing temperatures. Climate change also results in more extreme weather (e.g. increased frequency and severity of droughts, storms, and floods), melting sea ice and glaciers, rising sea levels, and much more.

2.2 Climate models and emission scenarios

A global climate model (GCM) is a computational model that simulates the climate system. They encode the physics and dynamics of the climate system via mathematical equations and are run on powerful supercomputers. Within the model, the Earth's atmosphere, ocean, and land is divided into a 3D grid of thousands of cells, with the size of the grid defining the resolution of the GCM (see schematic in Figure 1.). Current GCMs are typically run at spatial scales of about 100km due to their vast computational expense. GCMs can be used to make predictions of different climate variables and phenomena, much like having a synthetic Earth. They are validated over the past and are found to have good agreement with observations.




Figure 1. Schematic of a GCM. Image source: NOAA

GCMs are a critical tool for assessing and understanding not only how the climate has changed in the past, but what changes we can expect in the future. Many GCMs have been developed by different modeling centers around the world. They each run the same experiments: over the past (from 1850 – present) and for the future (present – 2100) and publish their data and results. In the Fifth phase of this Coupled Model Intercomparison Project (CMIP5) over 30 GCMs participated. Since the main driver of climate change is human activities, the GCMs run several “concentration trajectories”, also known as “emission scenarios”, in order to capture the range of possible future pathways that will vary according to human activity. These are known as Representative Concentration Pathways (RCPs). Two main pathways are often considered: RCP 4.5, which is an intermediate scenario (GHG emissions peak in 2040 and then decline) and RCP 8.5, which is a worst-case scenario (GHG emissions continue to rise through the end of the century). The Fifth Assessment of the United Nations Intergovernmental Panel on Climate Change (IPCC) report used RCPs 4.5 and 8.5 [IPCC, 2014]. The sixth phase of CMIP and the IPCC report (finalized in March 2023) [IPCC, 2022] added further refinement to the RCPs by using Shared Socioeconomic Pathways (SSPs), but those are not considered in this report.

It is important to note the differences between weather/climate and weather forecasts/climate projections. Weather refers to short-term conditions of the atmosphere on the timescales of days or weeks, whereas climate refers to long-term changes. The timescale of climate projections used in this report is daily, however these projections can’t be used in the same way as weather forecasts. It isn’t possible to make a prediction for a given day in the future. However, what they can be used for is to help us understand what to expect in the future *in a general sense*. In climate, we speak of trends and averages on periods of 30 years or longer. For example, the temperature trends for California are increasing, though there is significant daily, seasonal, even yearly variability due to the chaotic nature of the climate system. Climate projections can also tell us on average how much more frequent and severe heatwaves will be for example.

2.3 Climate impacts at a regional scale

The typical outputs retrieved from GCM simulations are often too coarse (~ 100km) in spatial resolution and/or too biased relative to observations to directly and reliably inform site specific infrastructure decisions. In order to assess climate impacts at a finer scale, it is necessary to downscale and bias correct the outputs from the GCMs. Multiple methodologies exist to downscale and refine the projections from GCMs to geographic and time scales appropriate for informing infrastructure planning decisions. The two traditional approaches to downscaling have been categorized as being “statistical” or “dynamical”. Novel approaches include “hybrid”, which combines statistical and dynamical aspects, and those that incorporate machine learning. Each approach has its own pros and cons. The statistical approach is often built from empirical statistical relationships between the simulated large-scale synoptic weather condition and the observed local conditions. In this report, we leverage the LOCA (Localized Constructed Analogs) statistical downscaling method, developed by Scripps Institution of Oceanography [Pierce et. al., 2014]. It uses historical observations to increase the resolution of the outputs from the GCMs. The resolution of the LOCA downscaled data that we will use for this report is 6km. The LOCA data we use was generated to support climate impact studies for the California 4th climate change

assessment.⁶ LOCA was used to downscale all 32 GCMs that contributed to CMIP5, though some of the downscaled climate variables are only available for a subset of 10 GCMs that we were identified as adequately sampling changes in California’s climate across the 32 GCMs. Of those 10, four GCMs were further selected that represent the range of projections from the 10 sub-selected GCMs. Considering results from an ensemble of GCMs ensures that we are capturing the range of possible outcomes due to the different ways each model represents the climate system. Ensuring that the GCMs that are used in impact assessments are producing accurate results over the historical period for the spatial location of interest is important. The sub-selected GCMs that were downscaled with LOCA and used in this report are found to accurately simulate the important aspects of the climate of California.

⁶ <https://www.climateassessment.ca.gov/>

3. Extreme heat and drought impacts to data centers

3.1 Assessing relevant impacts

Raw outputs from a GCM or from a subsequent downscaled method such as LOCA are often not in a form that is readily usable by an end-user. For example, having the temperature for a data center for every day until 2100 isn't useful by itself, since climate data isn't meaningful when looked at on a daily scale. Oftentimes we will want to perform some analysis on or synthesis of this data to obtain a more useful quantity for the end-user of the climate information. We'll refer to these derived quantities as climate analytics. An example of a climate analytic that we'll encounter in the coming chapter is the *average number of days in a year where the maximum daily temperature (a climate variable) exceeds a particular threshold*. The space of possible climate variables and analytics is huge, so a vital part of conducting a useful climate impacts assessment is to identify the relevant climate variables and analytics for the sector of interest. Through stakeholder engagement with several data centers in Northern California, we were able to provide a tailored assessment of the impacts of extreme heat and drought on data centers.

Not all data centers are equally at risk from extreme heat and drought. One key aspect that determines their vulnerability is the cooling methods they adopt. Air-cooled data centers are vulnerable to high *dry-bulb temperatures*. High temperatures can cause failure of the cooling systems if temperatures surpass the data center's design limits. Water cooling is typically more cost-effective and efficient (since water has a higher thermal conductivity than air) and therefore is a popular method for data center cooling, especially for high-power density computing. There is a trade-off: using water for cooling can reduce the amount of power a data center uses and so is more energy efficient, but instead they heavily depend on water, which is problematic in water scarce regions.

Many mid- to large-sized data centers use a chilled water system, which distributes cool water to the server room cooling units (see Figure 2.). Water is primarily consumed through evaporation from the cooling tower and through "blow-down". Blow-down is when the cooling tower dumps water to eliminate the buildup of contaminants that occurs after several cycles. The amount of blow-down varies as a function of the water quality and treatment. A typical large-sized water-cooled data center can consume up to 1-5 million gallons of water per day.⁷ This puts the data centers at high risk from water shortages that can arise during periods of *drought*.

⁷ <https://www.washingtonpost.com/climate-environment/2023/04/25/data-centers-drought-water-use/>




Figure 2: Schematic of a chilled water system. Image source: <https://dc.mynetworkinsights.com/data-center-cooling-infrastructure/>

The evaporation from the cooling towers is how the data center is cooled in this system. Therefore, water-cooled data centers are also vulnerable to high *humidity* or *wet-bulb temperatures*. Standard air temperature that people most often refer to is known as the dry-bulb temperature. Wet-bulb temperature accounts for humidity in the air and is the temperature of adiabatic saturation. It can be measured using a thermometer with the bulb wrapped in a wet cloth. When the relative humidity of the air is 100%, the water on the cloth is unable to evaporate and the wet-bulb temperature is the same as the dry-bulb temperature. However, when the humidity is lower, water from the wet cloth can evaporate and therefore the wet-bulb temperature can be lower than the dry-bulb temperature. This is like the effect of sweating as a means to cool down. If wet-bulb temperatures are higher than was assumed in the design of the data center, the water in the cooling towers cannot evaporate efficiently and the data center loses its cooling capability. Another common water-cooled method for data centers is Direct Evaporative Cooling (DEC), a method that cools outside air by using a wetted medium within an air handling unit. This method also is vulnerable to drought and wet-bulb temperature due to the same reasons above. Air-cooled data centers can also suffer as a result of high humidity since the air conditioning systems will have to work harder to remove the humidity from the outside air.

Increased *warm nights*, or minimum temperatures, can put data centers at risk as it can mean that the data centers cannot get relief at night from high temperatures. For example, a string of consecutive extreme heat days and warm nights could result in the temperature of the water used for cooling to be too high provide adequate cooling.

3.2 Data center relevant impacts

Extreme heat (maximum dry-bulb temperature)

In the Bay Area, annual maximum temperatures are expected to increase significantly. Figure 3. shows the time series and time period summaries of the annual average maximum temperatures for observations, and the 10 sub-selected CMIP5 GCMs downscaled with LOCA.

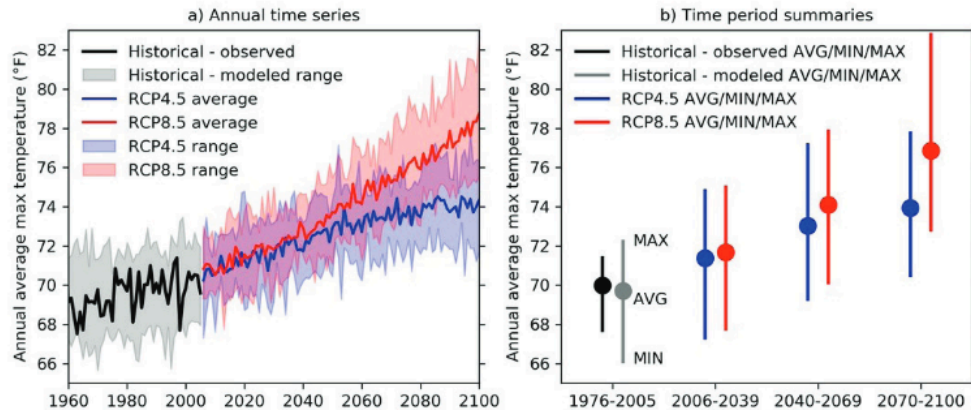


Figure 3: The time series and time period summaries of the annual average maximum temperatures for observations, and the 10 sub-selected CMIP5 GCMs downscaled with LOCA. The average across the 10 GCMs is shown for the RCP 4.5 and 8.5 scenarios and their ranges. Image source: https://www.energy.ca.gov/sites/default/files/2019-11/Reg_Report-SUM-CCCA4-2018-005_SanFranciscoBayArea_ADA.pdf

In Figure 4., we can see that the average hottest day of the is also expected to increase dramatically in the Bay Area region. The impacts of climate change on this region are already being felt, with record maximum temperatures in recent years.

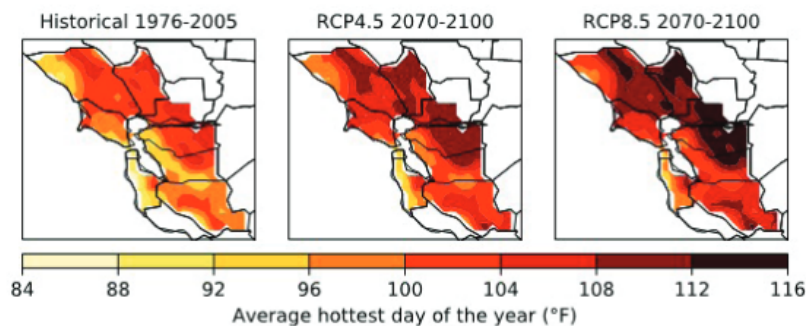


Figure 4: Average hottest day of the year across for historical, RCP 4.5 and RCP 8.5 future scenarios for the 10 sub-selected CMIP5 GCMs downscaled with LOCA. Image source: https://www.energy.ca.gov/sites/default/files/2019-11/Reg_Report-SUM-CCCA4-2018-005_SanFranciscoBayArea_ADA.pdf

Similar trends are seen for other parts of Northern California. Details for these other regions can be found in the California climate change assessment regional reports.⁸

⁸ <https://www.climateassessment.ca.gov/regions/>

Extreme heat analytics that we consider in the following chapter are as follows:

- Average number of extreme heat days per year above a user-defined threshold.
- Average number of heatwaves (a user-defined number of consecutive days) above a user-defined temperature.
- Average number of days in longest stretch of extreme heat days above a user-defined threshold.

Warm nights (minimum dry-bulb temperature)

The minimum temperature (which typically occurs during the night) is also projected to increase. Figure 5. shows the annual average minimum temperature under the RCP 8.5 scenario.




Figure 5: Time series of the annual average minimum temperature under the RCP 8.5 scenario for the 10 LOCA-downscaled GCMs for the Bay Area region, where the gray shading shows the projected range from all 32 GCMs. Image source: Cal-Adapt.org

The trends for other regions in Northern California can be found on the Cal-Adapt platform.

Like extreme heat, the analytics that we consider in the following chapter for warm nights are as follows:

- Average number of warm nights per year above a user-defined threshold.
- Average number of times the minimum temperature surpasses a threshold temperature for a user-defined number of consecutive days.
- Average number of days in the longest stretch of consecutive warm nights above a user-defined threshold.

Relative humidity and maximum wet-bulb temperature

Relative humidity is the ratio of the water vapor in the air to how much water vapor the air could potentially contain at its current temperature, expressed as a percentage. As the atmosphere warms, it can hold much more moisture. It is a typical output from a GCM. Wet-bulb temperature on other hand is not an output from GCMs, but it can be estimated empirically from the relative humidity and the dry-bulb temperature. In the next chapter, following the approach in Alessi et. al. [2020], our maximum wet-bulb temperature projections are calculated using daily maximum temperature and daily minimum relative humidity with the following empirical equation [Stull, 2011]:

$$T_w = T \operatorname{atan}[0.151977(\operatorname{RH}\% + 8.313659)^{1/2}] + \operatorname{atan}(T + \operatorname{RH}\%) - \operatorname{atan}(\operatorname{RH}\% - 1.676331) + 0.00391838(\operatorname{RH}\%)^{3/2} \operatorname{atan}(0.023101\operatorname{RH}\%) - 4.686035.$$

We won't consider it in this report, but the minimum wet-bulb temperature can be similarly calculated using the daily minimum dry-bulb temperature and daily maximum relative humidity.

We will explore the same analytics in the following chapter as above for minimum relative humidity, maximum relative humidity, and maximum wet-bulb temperature. We will also consider analytics of combinations of these quantities.

Drought

Data centers that use water for cooling are vulnerable to drought. Drought is challenging to both define and quantify. Generally, there four different types of drought: 1) meteorological, 2) hydrological, 3) agricultural, and 4) socioeconomic. Meteorological drought is often defined as a deficiency of precipitation over an extended period of time. However, drought severity and duration increases due to rising temperatures causing enhanced evaporation, even if there is no change to the amounts of precipitation [Wehner et al. 2017]. Hydrological drought cascades from meteorological drought, when the lack of precipitation impacts the water storage and supply, such as in streams, reservoirs, and groundwater levels, typically after a prolonged meteorological drought. It is this type of drought that is most relevant for data centers. From this definition, we can see that a meteorological drought in one region can cause a hydrological drought in another if the former region supplies water to the latter. This makes it critical for a data center to monitor meteorological drought conditions over the regions that supply their water. For example, a meteorological drought in Utah and Colorado can diminish water supplies in Lake Mead and Lake Powell, both of which are key reservoirs for California water supply.

The Southwestern US is expected to face significant drought risk in the coming decades [Cook et. al., 2015] Figure 6. shows the Palmer Drought Severity Index (PDSI; the most prominent meteorological drought index that is used widely for drought monitoring within the US), computed for 17 CMIP5 GCMs. All models are projecting unprecedented drought risk in the future for this region.

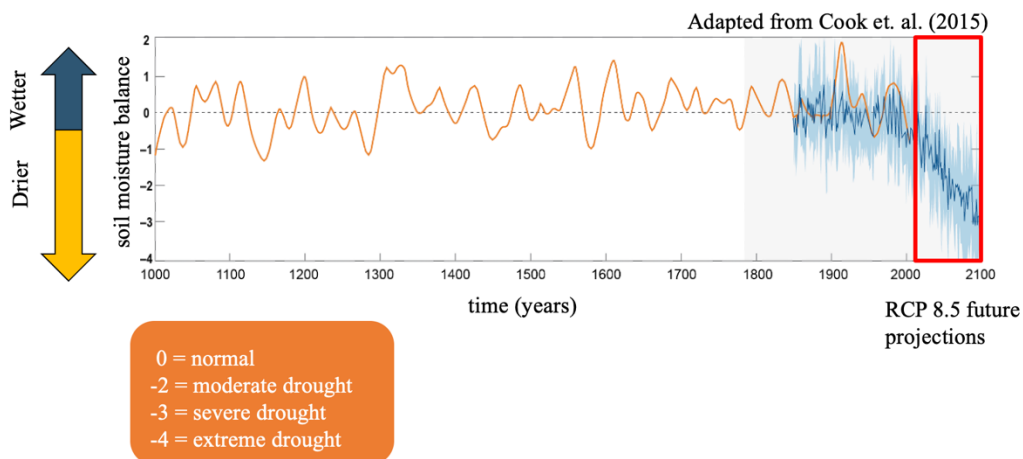


Figure 6: PDSI from observations (tree-ring data) in brown. Dark blue is the multi-GCM mean prediction for the RCP 8.5 future scenario computed for 17 CMIP5 GCMs. The light blue shaded area is the multi-GCM interquartile range showing the range of projected values. Image source: Cook et. al. 2015.

Sierra Nevada snowpack is a key water source that is increasingly at risk. In 2022, DWR reported that the snowpack was the fifth smallest on record since 1950, at a mere 35% of normal. Snow accumulated in the Sierra Nevada during the winter months is slowly released through the spring and summer months as snowmelt and then used as a primary water resource, supplying ~ 60% of Bay Area water.⁹ A recent study by Mote et al. [2018] found that average snowpack in the Western U.S. has declined 15-30% since 1915. A primary driver of the decline is the rising temperatures, a result of human induced climate change. This causes much of the winter precipitation to fall as rain and leads to earlier snowmelt, resulting in depleted water resources throughout the summer months [Fyfe et al., 2017, Kapnick and Hall, 2012, Pierce et al. 2008]. Furthermore, under the RCP 8.5 scenario, the average Sierra Nevada snowpack is projected to decline by nearly 83% by 2075-2100 [Rhoades et al., 2018]

By 2040, the Department of Water Resources estimates that California could lose 10% of its water supplies. Significant water shortages are to be expected in the Northern California region in the decades to come, putting data centers in this region at risk. Last year the California State Water Board adopted emergency water use regulations. To date, data centers have not been impacted by California’s water restrictions, but as water becomes a scarcer resource this may no longer be true in the future.

Droughts can also lead to increased energy costs as a result of declined hydroelectric availability. In a typical year, 15% of California’s electricity comes from hydroelectric power. Due to extreme drought, that number has fell by 48% below the 10-year average (2011–2020) in recent years.¹⁰

⁹ https://www.energy.ca.gov/sites/default/files/2019-11/Reg_Report-SUM-CCCA4-2018-005_SanFranciscoBayArea_ADA.pdf

¹⁰ <https://www.eia.gov/todayinenergy/detail.php?id=51839>

4. Case studies

In order to facilitate more interaction with the LOCA downscaled GCM data, we built a geospatial system. The system sends a json request to Cal-Adapt REST API requesting the data (e.g model, scenario, variable, time period.). See Figure 7. Our system in the backend will postprocess the data based on the filters and equations we provide.



Figure 7: How our geospatial system fetches data

The system was developed with Bokeh¹¹ and all of the data that we use is publicly available via the Cal-Adapt API. It is worth noting that some of the results that we can obtain with our system can also be obtained via Cal-Adapt’s platform. We go beyond the platform in some ways that make sense for data center-specific impacts. For example, we can look at combinations of variables such as extreme heat and humidity, and we have an interactive map with all the known data centers in Northern California. We’ll show some results for three locations as case studies. To avoid focusing on specific data centers, we’ll show the results for San Jose, Livermore, and Napa airports.

4.1 San Jose airport

First for San Jose, we’ll look at the maximum dry-bulb temperatures in Figure 8. We set it to the RCP 8.5 scenario and select all four GCMs. We set the threshold temperature to 100F and plot the number of days a year above 100F for each model. In the “Avg” row, we calculate the average number of days a year above 100F averaged across that time period and the four GCMs in this case. We can also look at one or a subset of the 4 GCMs if desired. The plot shows the average number of days above 100F through the end of the century, however in the “Avg” row, we adjusted the slider so that a 30-year time period was selected from 2025-2055. As previously mentioned, climate projection cannot tell us what temperatures to expect in a given year, but averages can be extremely insightful. 30 years is a reasonable estimate of the life-cycle of a data center. We can see that if we were to design our data center from historical observations alone, we would deduce that the average number of days above 100F for the historical period would be 0 and the maximum dry-bulb temperature to ever have occurred was 108F. We see that in the future projections that the average number of days a year above 100F, even just for the next 30 years, is three, and the maximum dry bulb temperature that is expected is 111F. These higher temperatures are of a concern not only since the data center infrastructure may not have been designed to withstand them, but in the case of water-cooled data centers higher temperatures will

¹¹ <https://bokeh.org/>

also increase the amount of water required for cooling. This could be problematic in drought years when water is limited, further amplifying the impacts of extreme heat to the data center.

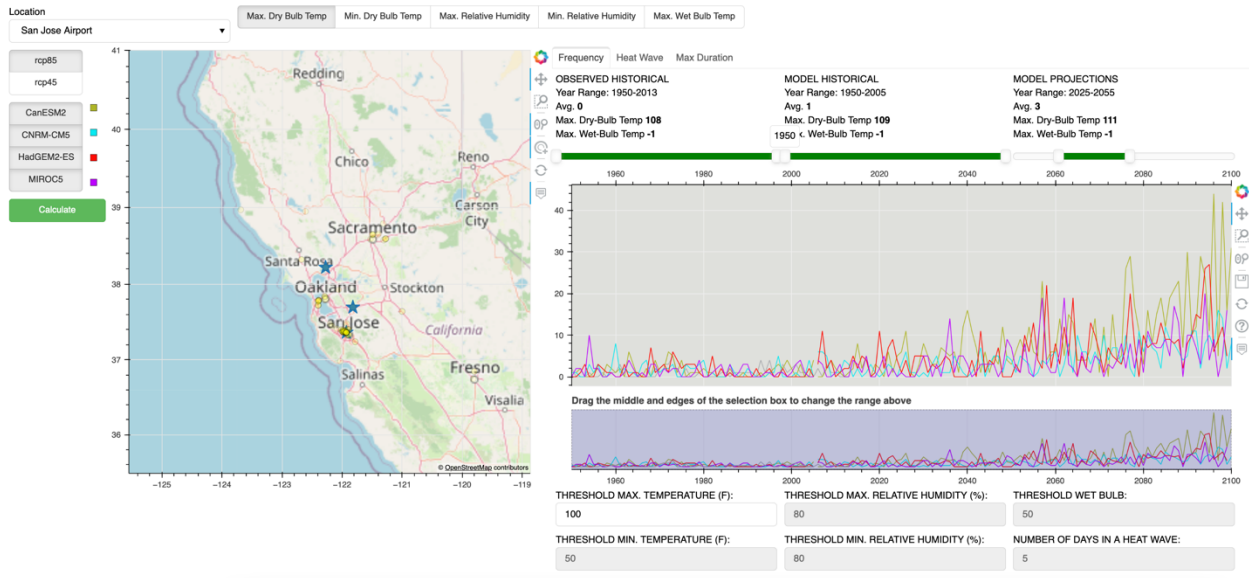


Figure 8: Dry-bulb temperatures for San Jose Airport

Looking at both extreme heat and maximum relative humidity together (i.e. occurring on the same day), with thresholds of 100F dry-bulb temperature and 80% relative humidity, we see that incidences of high heat and humid days is expected to increase. In Figure 9. we show the results for one GCM.

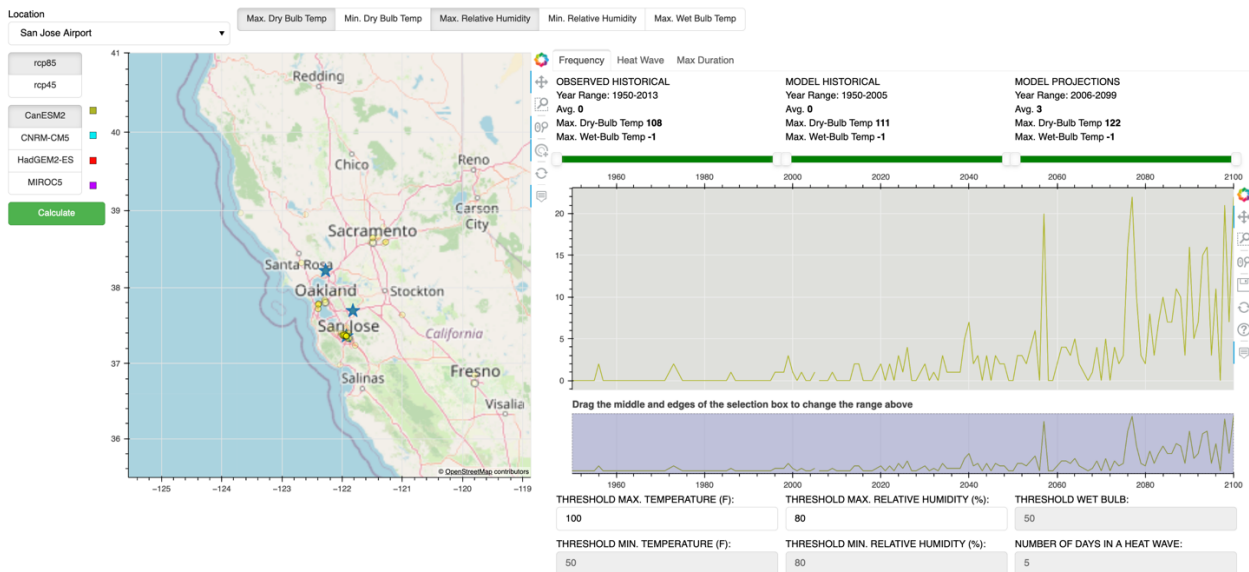


Figure 9: Combined impact of maximum relative humidity and dry-bulb temperature at San Jose Airport

4.2 Livermore airport

For Livermore we'll look at maximum wet-bulb temperatures above 78F in Figure 10. We see that within the next ~ 30 years the wet-bulb temperatures will exceed 78F, reaching maximums of 79F in that timeframe. One of the models is predicting there to be one year within the next ~ 30 years that has 12 days above 78F.

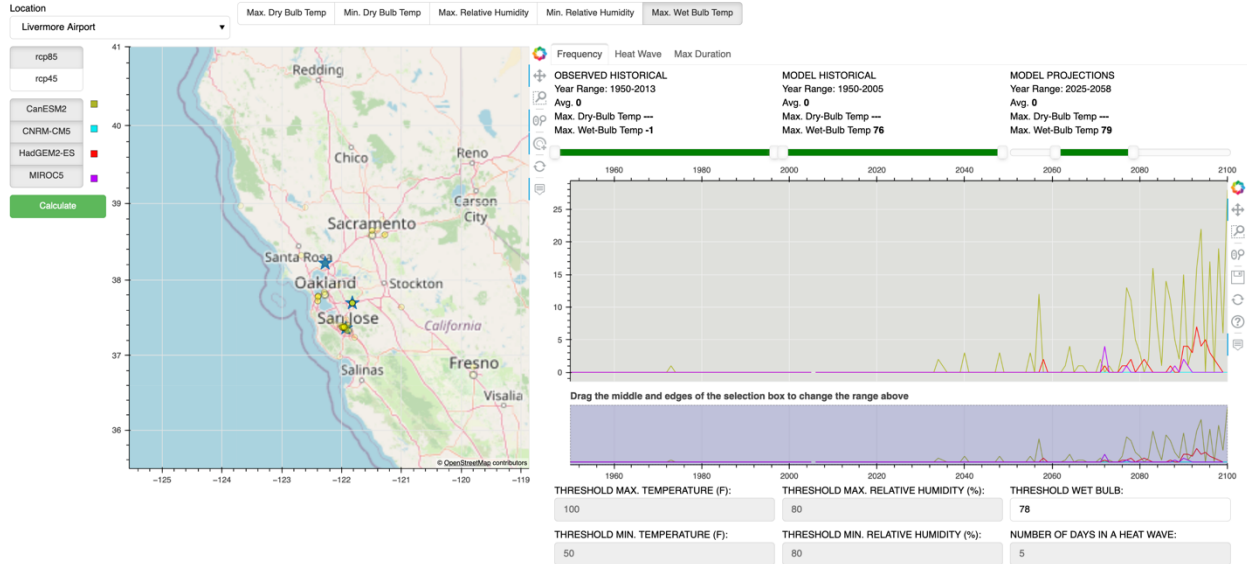


Figure 10: Maximum wet-bulb temperatures for Livermore Airport

4.3 Napa airport

For Napa airport we'll look at heatwaves in Figure 11. We specify the threshold temperature to be 100F and the heatwave length, i.e. the number of consecutive days above that threshold temperature to be 5. There were very few events of this nature historically, but in the future they are increasing significantly in frequency.




Figure 11: Heatwaves for Napa Airport

Looking a bit further out to mid-late century, in Figure 12. we can also see increases in the incidence of high dry-bulb and wet-bulb temperatures co-occurring.




Figure 12: Maximum dry and wet bulb temperatures for Napa Airport

5. The path forward

Data centers that use historical weather data such as Typical Meteorological Years (TMY) in their design and operations are putting themselves at risk from the impacts of climate change. Extreme heat and drought are expected to increase in severity, frequency, and duration in Northern California in the future, resulting in conditions that the data centers were potentially not designed and built to withstand. If action is not taken to incorporate future climate information into data center design and operations, outages, like those faced by Google and Oracle in the UK last year, will become much more common and widespread. Analyzing the climate variables and analytics described in this report is a necessary first step in developing a suitable resiliency strategy. Data centers owner and operators or other key stakeholders can use the LOCA downscaled data and Cal-Adapt platform to conduct analyses similar to those described in this report to determine their risk from extreme heat and humidity for locations in California. As well as ensuring that their data center can withstand the maximum temperatures and humidity expected in the future, they can also explore how many days on average per year a temperature is *below* a certain threshold. This can further assist with planning which cooling technologies to adopt. This analysis can also be used to help ensure that adequate redundancies and back-ups are considered. It is also vitally important for water-cooled data centers to stay informed of the region's water situation and be proactive to mitigate potential risks.

In California, there may be issues in the near future for data centers looking to obtain insurance that covers climate-related risk. For example, this year State Farm has ceased accepting applications for most types of new insurance policies in the state because of “rapidly growing catastrophe exposure.”¹² Allstate also recently stopped offering insurance to homeowners in California, and other states are also experiencing a similar trend.¹³ Therefore, taking a proactive approach to climate-related risks is prudent.

Data centers are under growing pressure to improve their sustainability, particularly reducing their vast power and water consumption. Implementing sustainable measures in data centers also has the added benefit of reducing their vulnerability to climate-related risks. For example, they can implement water-efficient cooling technologies such as air-side economizers or closed-loop cooling systems that will reduce their water consumption. Using water-recycling systems or alternative cooling methods like liquid immersion cooling can further minimize water reliance. This has the benefit of meeting sustainability goals and makes data centers less susceptible to water scarcity during a drought. Hyperscalers are already moving towards innovative cooling technologies in an effort to curb their water usage. For example, by 2030 Amazon Web Services (AWS) and Facebook plan to be water-positive.

A word on power. Due to threats of wildfire in extremely hot, dry, and windy conditions, California implements Public Safety Power Shutoffs (PSPS) that sees utilities turning off electricity if there is a threat to a portion of the electric system. Transitioning to renewable energy sources both reduces the carbon footprint of data center and has the benefit of reducing their reliance on the

¹² <https://newsroom.statefarm.com/state-farm-general-insurance-company-california-new-business-update/>

¹³ <https://www.usatoday.com/story/news/nation/2023/06/11/climate-change-effects-hit-us-homeowner-insurance/70288893007/>

traditional power grid which is vulnerable to climate-related hazards, including extreme heat and drought.

Computational fluid dynamic (CFD) simulations are a common tool typically employed by data centers to monitor and optimize the internal data center environment and ensure that the computer systems stay within optimal operational temperatures. They can be helpful when adding equipment, since the addition of the new equipment can be first modeled by the CFD simulation. CFD simulation can also be leveraged for external modeling as a tool to design optimal cooling systems.¹⁴ To date, as far as the author is aware, this has not been performed using future climate data. This could be a valuable approach to developing and implementing cooling systems that can withstand the climate of the future.

¹⁴ <https://www.futurefacilities.com/uploads/media/casestudy-ka0.pdf>

References

Alessi, M. J. and DeGaetano, A. T., 2020: Future Extreme Hourly Wet Bulb Temperatures Using Downscaled Climate Model Projections of Temperature and Relative Humidity. *Theor Appl Climatol*. doi: 10.1007/s00704-020-03368-0.

Arias, P.A., N. Bellouin, E. Coppola, R.G. Jones, G. Krinner, J. Marotzke, V. Naik, M.D. Palmer, G.-K. Plattner, J. Rogelj, M. Rojas, J. Sillmann, T. Storelvmo, P.W. Thorne, B. Trewin, K. Achuta Rao, B. Adhikary, R.P. Allan, K. Armour, G. Bala, R. Barimalala, S. Berger, J.G. Canadell, C. Cassou, A. Cherchi, W. Collins, W.D. Collins, S.L. Connors, S. Corti, F. Cruz, F.J. Dentener, C. Dereczynski, A. Di Luca, A. Diongue Niang, F.J. Doblas-Reyes, A. Dosio, H. Douville, F. Engelbrecht, V. Eyring, E. Fischer, P. Forster, B. Fox-Kemper, J.S. Fuglestedt, J.C. Fyfe, N.P. Gillett, L. Goldfarb, I. Gorodetskaya, J.M. Gutierrez, R. Hamdi, E. Hawkins, H.T. Hewitt, P. Hope, A.S. Islam, C. Jones, D.S. Kaufman, R.E. Kopp, Y. Kosaka, J. Kossin, S. Krakovska, J.-Y. Lee, J. Li, T. Mauritsen, T.K. Maycock, M. Meinshausen, S.-K. Min, P.M.S. Monteiro, T. Ngo-Duc, F. Otto, I. Pinto, A. Pirani, K. Raghavan, R. Ranasinghe, A.C. Ruane, L. Ruiz, J.-B. Sallée, B.H. Samset, S. Sathyendranath, S.I. Seneviratne, A.A. Sörensson, S. Szopa, I. Takayabu, A.-M. Tréguier, B. van den Hurk, R. Vautard, K. von Schuckmann, S. Zaehle, X. Zhang, and K. Zickfeld, 2021: Technical Summary. In *Climate Change 2021: The Physical Science Basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change* [Masson-Delmotte, V., P. Zhai, A. Pirani, S.L. Connors, C. Péan, S. Berger, N. Caud, Y. Chen, L. Goldfarb, M.I. Gomis, M. Huang, K. Leitzell, E. Lonnoy, J.B.R. Matthews, T.K. Maycock, T. Waterfield, O. Yelekçi, R. Yu, and B. Zhou (eds.)]. Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA, pp. 33–144. doi:10.1017/9781009157896.002.

Cook, Benjamin et al., Unprecedented 21st century drought risk in the American Southwest and Central Plains. *Sci. Adv.* 1,e1400082(2015).DOI:10.1126/sciadv.1400082

Fyfe, John C., et al. "Large near-term projected snowpack loss over the western United States." *Nature communications* 8 (2017): 14996.

IPCC, 2014: *Climate Change 2014: Synthesis Report. Contribution of Working Groups I, II and III to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change* [Core Writing Team, R.K. Pachauri and L.A. Meyer (eds.)]. IPCC, Geneva, Switzerland, 151 pp.

IPCC, 2022: *Climate Change 2022: Impacts, Adaptation, and Vulnerability. Contribution of Working Group II to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change* [H.-O. Pörtner, D.C. Roberts, M. Tignor, E.S. Poloczanska, K. Mintenbeck, A. Alegria, M. Craig, S. Langsdorf, S. Löschke, V. Möller, A. Okem, B. Rama (eds.)]. Cambridge University Press, Cambridge, UK and New York, NY, USA, 3056 pp., doi:10.1017/9781009325844.

Kapnick, Sarah, and Alex Hall. "Causes of recent changes in western North American snowpack." *Climate Dynamics* 38.9-10 (2012): 1885-1899.

Mote, Philip W., et al. "Dramatic declines in snowpack in the western US." *npj Climate and Atmospheric Science* 1.1 (2018): 2.

Pierce, David W., et al. "Attribution of declining western US snowpack to human effects." *Journal of Climate* 21.23 (2008): 6425-6444.

Pierce, D. W., Cayan, D. R., & Thrasher, B. L. (2014). Statistical downscaling using localized constructed analogs (LOCA). *Journal of Hydrometeorology*, 15(6), 2558– 2585. <https://doi.org/10.1175/jhm-d-14-0082.1>

Rhoades, Alan M., Paul A. Ullrich, and Colin M. Zarzycki. "Projecting 21st century snowpack trends in western USA mountains using variable-resolution CESM." *Climate Dynamics* 50.1-2 (2018): 261-288.

Stull, Roland. "Wet-bulb temperature from relative humidity and air temperature." *Journal of applied meteorology and climatology* 50.11 (2011): 2267-2269.

Wehner, M.F., J.R. Arnold, T. Knutson, K.E. Kunkel, and A.N. LeGrande, 2017: Droughts, floods, and wildfires. In: Climate Science Special Report: Fourth National Climate Assessment, Volume I [Wuebbles, D.J., D.W. Fahey, K.A. Hibbard, D.J. Dokken, B.C. Stewart, and T.K. Maycock (eds.)]. U.S. Global Change Research Program, Washington, DC, USA, pp. 231-256 doi: 10.7930/J0CJ8BNN

Williams, A.P., Cook, B.I. & Smerdon, J.E. Rapid intensification of the emerging southwestern North American megadrought in 2020–2021. *Nat. Clim. Chang.* 12, 232–234 (2022). <https://doi.org/10.1038/s41558-022-01290-z>