Pacific Northwest
NATIONAL LABORATORY

PNNL-34418

Building ControlScore: General Service Administration Office Building Deployment

June 2023

Tim A Yoder
Bryan C Pamintuan
Timothy I Salsbury

U.S. DEPARTMENT OF ENERGY

Prepared for the U.S. Department of Energy
under Contract DE-AC05-76RL01830

PACIFIC NORTHWEST NATIONAL LABORATORY
operated by
BATTELLE
for the
UNITED STATES DEPARTMENT OF ENERGY
under Contract DE-AC05-76RL01830

Printed in the United States of America

Available to DOE and DOE contractors from
the Office of Scientific and Technical Information,
P.O. Box 62, Oak Ridge, TN 37831-0062
www.osti.gov
ph: (865) 576-8401
fox: (865) 576-5728
email: reports@osti.gov

Available to the public from the National Technical Information Service
5301 Shawnee Rd., Alexandria, VA 22312
ph: (800) 553-NTIS (6847)
or (703) 605-6000
email: info@ntis.gov
Online ordering: http://www.ntis.gov

PNNL-34418

Building ControlScore: General Service Administration Office Building Deployment

June 2023

Tim A Yoder
Bryan C Pamintuan
Timothy I Salsbury

Prepared for
the U.S. Department of Energy
under Contract DE-AC05-76RL01830

Pacific Northwest National Laboratory
Richland, Washington 99354

Summary

Improvements to building control systems can lead to energy savings and increased occupant comfort. In an optimized system, process variables such as air temperature will closely follow their desired setpoints and avoid excess energy use. Typically, experts must manually inspect individual control loops to identify poor performance and opportunities for improvement. However, this approach is difficult in modern buildings that have a prohibitively large number of controllers. To address this issue, Pacific Northwest National Laboratory (PNNL) created the ControlScore concept, which takes operating data from the many controllers within a building and generates standardized scores for each loop on a scale of 0 to 10 (a score of 0 indicates poor control, a score of 10 indicates good control).

PNNL applied the Building ControlScore application to all available data from a General Services Administration office building within the period of January 1, 2023, to March 9, 2023. The building scored a 4.7 overall, with all 74 of the building's loops fitting a roughly normal distribution centered around 5. These results indicate that the analyzed systems have below-average performance with room for improvement, especially in the poorly scored systems. Airflow loops tended to have much lower scores than zone temperature loops. The lowest and highest performing systems in the building section were identified, as were all loops with a score less than 1.

While the ControlScore identifies loops and systems that aren't meeting their designated setpoints, it does not indicate the cause of those issues. For example, consider a supply air terminal unit's airflow loop that received a low score due to it delivering less air than specified by the setpoint. The lower-than-desired airflow could be due to equipment limitations (e.g., the terminal unit or duct serving is too small to accommodate that airflow), malfunctioning equipment (e.g., a stuck damper or bad sensor), or something else entirely. The ControlScore does not diagnose problems it simply identifies the symptoms that can be explored and addressed by building operators.

Acknowledgments

We would like to thank the General Service Administration's Center for Emerging Building Technologies, property management, and building management specialists for their support in providing access to the building data and feedback on the tool design.

Acronyms and Abbreviations

GSA	General Services Administration
ID	identification
PNNL	Pacific Northwest National Laboratory
PV	process variable
SP	set point
VAV	variable air volume

Contents

Figures

Tables

1.0 Introduction

This report provides an overview of Pacific Northwest National Laboratory's (PNNL's) Building ControlScore application, including the methodology behind ControlScore, and the results from deployment of the regulatory control layer of ControlScore which includes the proximity and stability aspects of the tool applied to a data set from a General Services Administration (GSA) office building.

1.1 Background

Modern buildings contain large numbers of feedback controllers that regulate process variables (PV) to follow the associated setpoints (SP). In this context, a PV is the thing being controlled (e.g., the actual zone temperature), and the SP is the desired target of that PV (e.g., a zone temperature setpoint). The performance of these controllers has a direct impact on occupant comfort and the overall energy use of a building. Thus, identifying poorly performing loops and systems is the first step to improving a building's control, comfort, and potentially energy efficiency metrics such as energy use intensity. Some buildings may use alarms to flag variables that cross a certain threshold, but this method is not as reliable as more sophisticated methods, such as the ControlScore's approach described in this report, that analyze behavior over time and generally produce a large number of "nuisance" alarms that operators frequently ignore.

1.2 ControlScore

Typically, investigating control performance requires a manual inspection of a system. Within the system of interest, the loops must be identified. For example, in a variable air volume system, each air terminal unit can have multiple control loops for airflow volume, discharge air temperature, and zone temperature. Then, the values of the PV and SP over time within that loop must be analyzed (e.g., comparing the trended histories of the actual airflow and the desired airflow). Determining the loop's control performance is not an exact and simple process. Rather, an expert must make the determination based on knowledge of the type of system. This approach may work for a few systems that have been flagged by alarms or can be identified by outlier behavior compared to similar systems in the same building, but this does not identify all the problematic loops, nor can the results be aggregated for comparison between buildings or sites. Inspection by an expert likely cannot be used for every single system in a building since modern buildings can have too many of these loops to monitor and evaluate the performance of each manually. Thus, there is a need for a more systematic, automated, and standardized method of scoring building control performance that can be applied to large datasets while avoiding prohibitively large workloads.

To address these problems, researchers at PNNL created the ControlScore concept (Salsbury 2023). ControlScore takes operating data from the many controllers within a building and generates standardized scores for each loop on a scale of 0 to 10 (a score of 0 indicates poor control, a score of 10 indicates good control) using a probability-based measure. The scoring methodology includes three aspects of control: proximity of PV to SP, stability of PV, and comparison of SP trajectory with references.

The proximity aspect measures how closely the PV is being controlled by the SP. The distribution of points in the PV is compared to the distribution of points in the SP. Figure 1 below illustrates two examples of this scoring method. The top image of Figure 1 shows an example of a PV that generally has the same value as the SP, so it receives a high score. In contrast, the bottom image in Figure 1 shows an example of a PV that is not close to the SP, so it receives a low score.

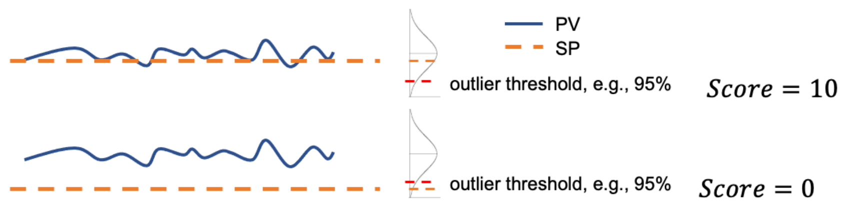

Figure 1. Proximity Aspect of ControlScore

The stability aspect measures how the PV fluctuates around the SP over time. Figure 2 below illustrates this concept, where the oscillations in PV can be scored. This behavior is not captured in the proximity score because the average error of both signals in Figure 2 may be close to 0, but the behavior represented in the top part of the figure may indicate poor control performance.

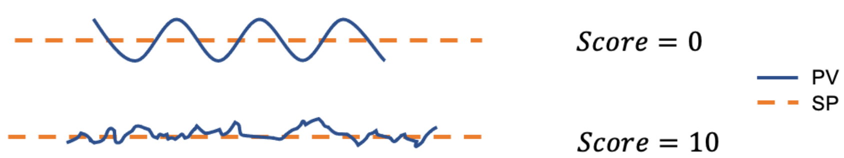

Figure 2. Stability Aspect of ControlScore

The expected trajectory aspect compares the actual SP to a reference SP that is known to be well-performing. Figure 3 illustrates this concept, where the score is a function of the distance between the actual SP and the well-performing reference.

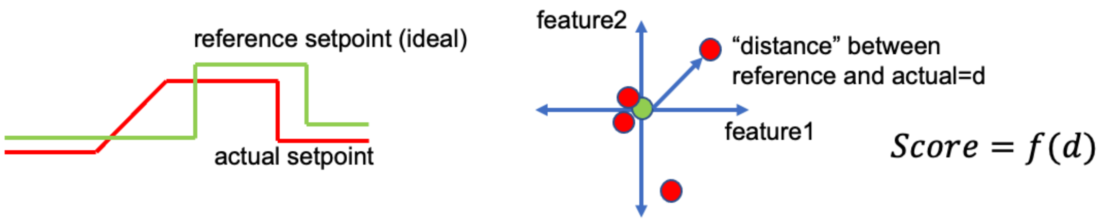

Figure 3. Trajectory Aspect of ControlScore

The regulatory control layer, which includes only the proximity and stability aspects, is the focus of this deployment of ControlScore. Supervisory control, which is the SP reference aspect, is not included in this analysis due to insufficient metadata for the loops and systems being assessed. Scores are calculated for each loop and then aggregated to the system, building, and site levels to allow for easy assessment. Calculated scores are stored in a database and accessible through a web-based user interface.

Calculating control scores requires certain data properties. The **data required** to produce a regulatory score is only the time series measurements (or trend history) of SP and the associated PV. Evaluation of the SP trajectory requires additional information. **Sample frequency** does not have a significant effect on the scoring methods as long as the dynamics of the loop are captured. However, very slow sampling may reduce the information contained in the signals and mask certain aspects of behavior, while very fast sampling may add noise. The **time span** should be monthly to both include enough samples while still showing changes over time. **Data quality** does not have a significant effect on the scoring methods, as long as there are not too many gaps or severe quantization.

2.0 Analysis

PNNL applied the Building ControlScore application to all available data from a GSA office building. This analysis covers two aspects of the ControlScore: the proximity of the PV to the SP and stability of the PV. The data covers January 1, 2023, to March 9, 2023.

2.1 Data Available and Metadata Extraction

For this analysis, the available data was manually extracted from Tridium Niagara user interface as we were unable to access the historian database. This manual effort involved exporting the histories of a single point at a time and only included data from the recent past (e.g., three months prior to the date the data was downloaded). Unfortunately, this process limited the extent of the systems the ControlScore tool could be applied to, as the process was prohibitively time consuming. Additionally, this deployment only covers variable air volume (VAV) terminal units, as they were the only systems that had the point trend histories set up for both the PV and SP points. Incorporating the air handling unit systems into the ControlScore tool would provide valuable insights on those systems, which have a large impact on the energy performance of the buildings, but we were unable to do so, as there were not SPs trended in the Niagara control system that we were able to access.

Additional metadata on the points, their relation to physical systems, how they related to each other, etc. had to be inferred from the data in the database.

The original data export included 38 different csv files, each containing the time series data of airflow (cfm), zone temperature (°F), and the associated setpoints for 1-3 VAVs. Each column (representing the values over time of a single variable) was separated into its own individual csv file along with the associated timestamps. Also, a list was compiled of each point’s name.

Each of these points can be categorized as a PV, SP, or a measurement point that is irrelevant to ControlScore (although all points in this data set were PV or SP). This point mapping allows PVs to be matched to SPs for scoring. These points can also be grouped into their respective loops and systems, which are used for proper matching of PVs to the correct SPs and allow aggregation of results. Ideally, this information would be modeled directly with some type of semantic model, such as one based on BrickSchema or Haystack. However, that type of metadata was not available for this analysis, so it needed to be extracted from the available data.

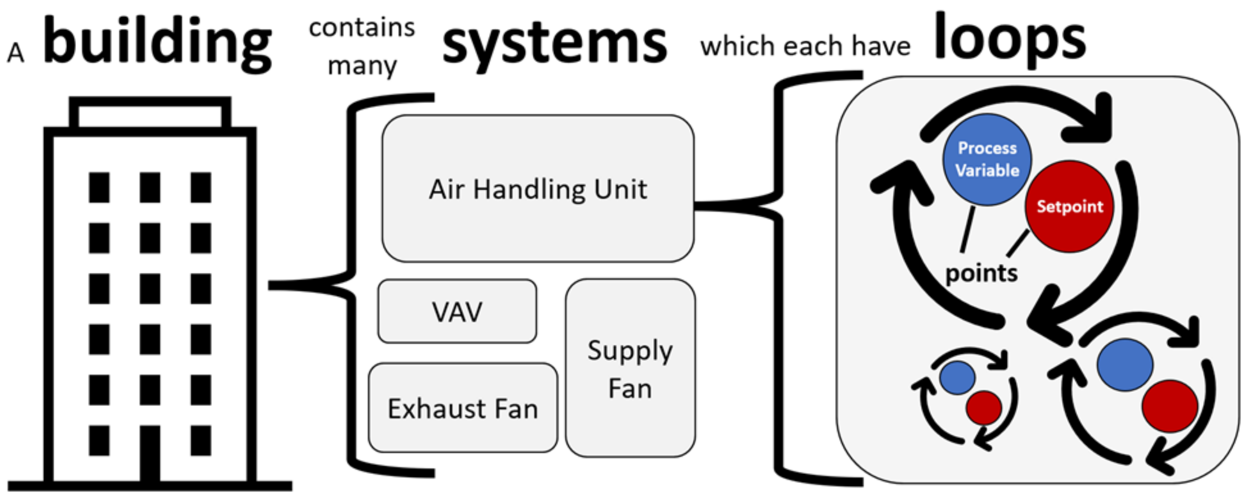

Figure 4. Hierarchy of Points, Loops, Systems, and Buildings

To do this point mapping, the necessary information was found in each point name in the compiled list. Splitting up a point name resulted in a key piece of information: the system name. The same information was extracted from each of the available points. An example summary for a small subset of points is shown below in the first two columns of Table 1.

Table 1. Information Extracted from Point Name

Full Point Name	System	Loop Name	Loop Type
VAV_1_Airflow_Setpoint	VAV_1	VAV_1_Airflow	Airflow
VAV_1_Control_Setpoint	VAV_1	VAV_1_Zone_Temperature	Zone_Temperature
VAV_1_Zone_Temperature	VAV_1	VAV_1_Zone_Temperature	Zone_Temperature
VAV_1_airflow	VAV_1	VAV_1_Airflow	Airflow
VAV_2_Airflow_Setpoint	VAV_2	VAV_2_Airflow	Airflow
VAV_2_Control_Setpoint	VAV_2	VAV_2_Zone_Temperature	Zone_Temperature
VAV_2_Zone_Temperature	VAV_2	VAV_2_Zone_Temperature	Zone_Temperature
VAV_2_airflow	VAV_2	VAV_2_Airflow	Airflow

VAV = variable air volume.

The system names were used to group points in the same system. Then, some manual inspection was done to assign loop names and loop types. Table 1 above shows points from two different VAV systems that each contained airflow and zone temperature loops. Next, each point was mapped to a point type (PV or SP). An example summary is shown below in Table 2.

Table 2. Point Mapping

Full Point Name	PV	SP
VAV_1_Airflow_Setpoint	0	1
VAV_1_Control_Setpoint	0	1
VAV1_Zone_Temperature	1	0
VAV_1_airflow	1	0
VAV_2_Airflow_Setpoint	0	1

VAV_2_Control_Setpoint	0	1
VAV_2_Zone_Temperature	1	0
VAV_2_airflow	1	0

PV = process variable; SP = set point;
VAV = variable air volume

Next, identifications (IDs) were assigned to each unique system, loop, and point. An example subset of the point ID table is shown below in Table 3. Similar tables exist to identify the system and loop. The IDs from these tables are also shown in the point ID table below.

Table 3. Point ID Table

Full Point Name	Point ID	Loop ID	System ID	PV	SP
VAV_1_Airflow_Setpoint	1	1	1	0	1
VAV_1_Control_Setpoint	2	2	1	0	1
VAV_1_Zone_Temperature	3	2	1	1	0
VAV_1_airflow	4	1	1	1	0
VAV_2_Airflow_Setpoint	5	3	2	0	1
VAV_2_Control_Setpoint	6	4	2	0	1
VAV_2_Zone_Temperature	7	4	2	1	0
VAV_2_airflow	8	3	2	1	0

ID = identification; PV = process variable; SP = set point; VAV = variable air volume.

The information from the ID tables allows each datapoint to be matched to its setpoint within the same loop and categorized into its appropriate system. This organization allows ControlScore to identify which PV should be regulated by which SP. In Table 3 above, points 1 and 4 (Point ID) are seen to be in the same loop (Loop ID 1); one is the PV and one is the SP. Similarly, points 6 and 7 are PV and SP pairs in a different loop (Loop ID 4), which is part of a different system. ControlScore was able to take this metadata and compare the appropriate PV and SP pairs to generate scores (according to the methodology described in Section 1.2), then aggregate the results at a system level.

2.2 Building Performance

The next few sections of this report will provide an overview of the results, trends, and a few examples of specific systems and loops with both high and low performance.

Analysis of the building data resulted in an overall building control score of 4.7, which is made up of 74 individual system scores. The trend of this overall score over time is shown in Figure 5 below. The behavior has remained relatively stable over time, with the score remaining between 4.5 and 5.5.

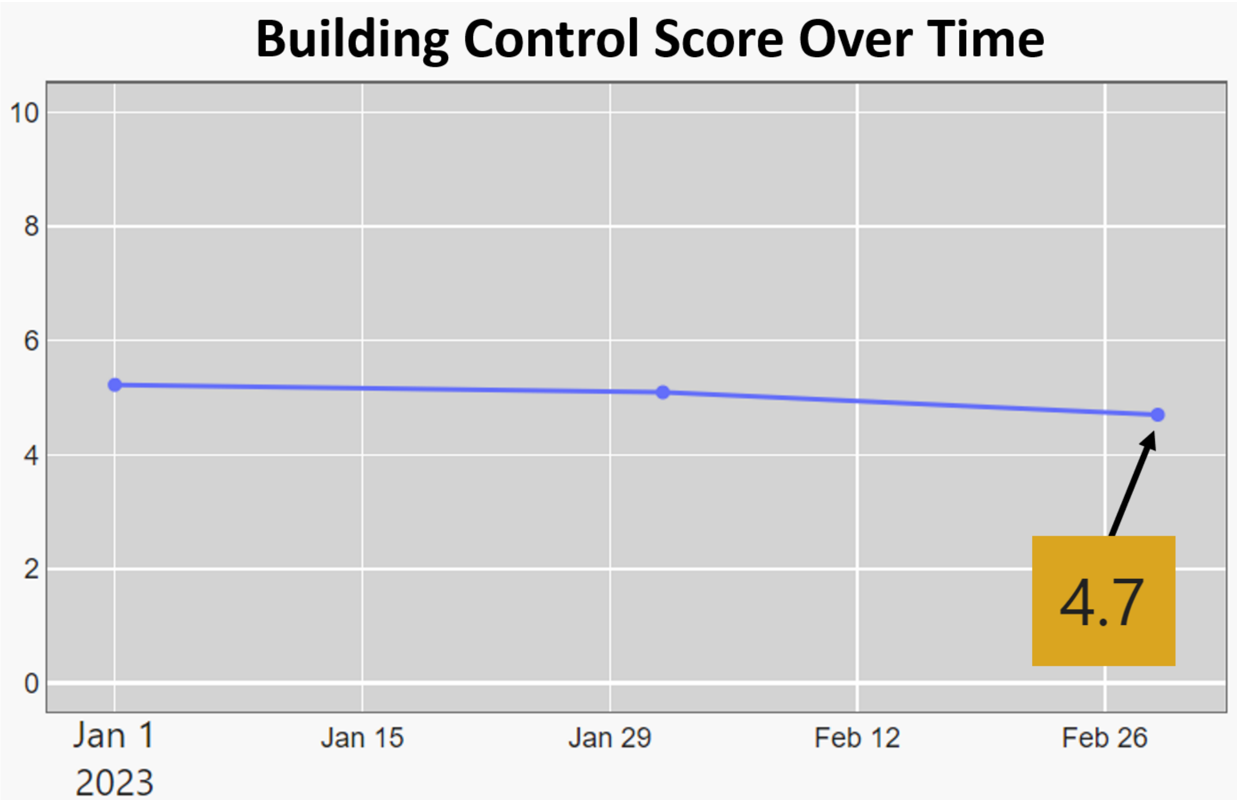

Figure 5. Building Level ControlScore Over Time

The GSA office building includes 74 different VAV systems, each with an airflow and zone temperature loop. The distribution for scores of these systems is shown in Figure 6 below. As can be seen, the scores are somewhat normally distributed. Further examination of the building equipment may reveal some common characteristics of poorly performing systems (e.g., older equipment).

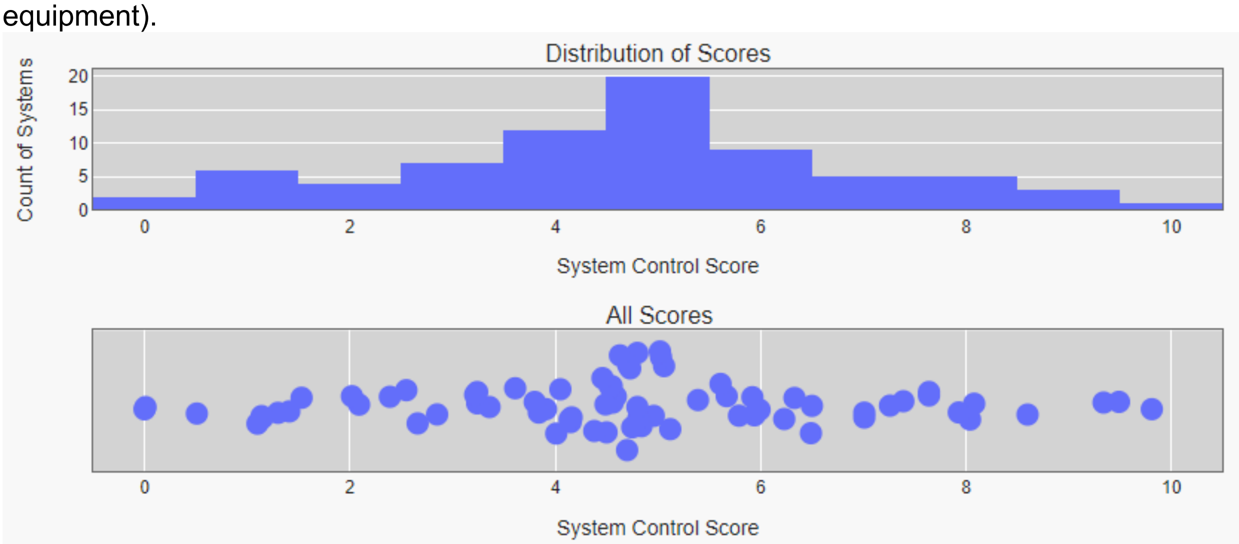

Figure 6. Distribution of 74 System Scores

These results indicate that the building has below-average overall performance with room for improvement, especially in poorly performing systems. Since the system scores are normally distributed, the poorly performing equipment should be examined first to find and resolve the issues causing low scores.

Table 4 below shows the systems with the highest 10 scores. The maximum system score achieved is 9.81.

Table 4. 10 Highest Scoring Systems

System ID	System Control Score	Airflow Score	Zone Temp Score
74	9.81	9.70	9.91
75	9.49	9.78	9.19
3	9.34	9.67	9.02
67	8.6	9.46	7.74
66	8.08	8.33	7.84
28	8.04	9.07	7.01
68	7.93	9.59	6.26
18	7.64	9.03	6.24
23	7.64	7.74	7.53
38	7.39	7.35	7.42

ID = identification; VAV = variable air volume.

Table 5 below shows the systems with the lowest 10 scores. Only one system has a minimum score of 0.

Table 5. 10 Lowest Scoring Systems

System ID	System Control Score	Airflow Score	Zone Temp Score
34	0.00	0.00	0.00
19	0.01	0.00	0.02
13	0.51	0.00	1.02
52	1.10	0.00	2.20
1	1.14	0.16	2.11
24	1.15	0.75	1.55
40	1.30	1.85	0.74
12	1.41	0.27	2.55
15	1.53	0.00	3.07
57	2.02	0.00	4.04

ID = identification; VAV = variable air volume

2.3 Airflow Loops

Figure 7 below shows the score over time and the raw data for the airflow loop in system 74. The control score over time shows an increasing trend from near 1 up to 9.7. The bottom left figure shows the raw time series data with the airflow setpoint in blue and the airflow process variable in red. The air flow follows the setpoint very tightly through the changing values without much oscillation, resulting in a high 9.7 loop control score.

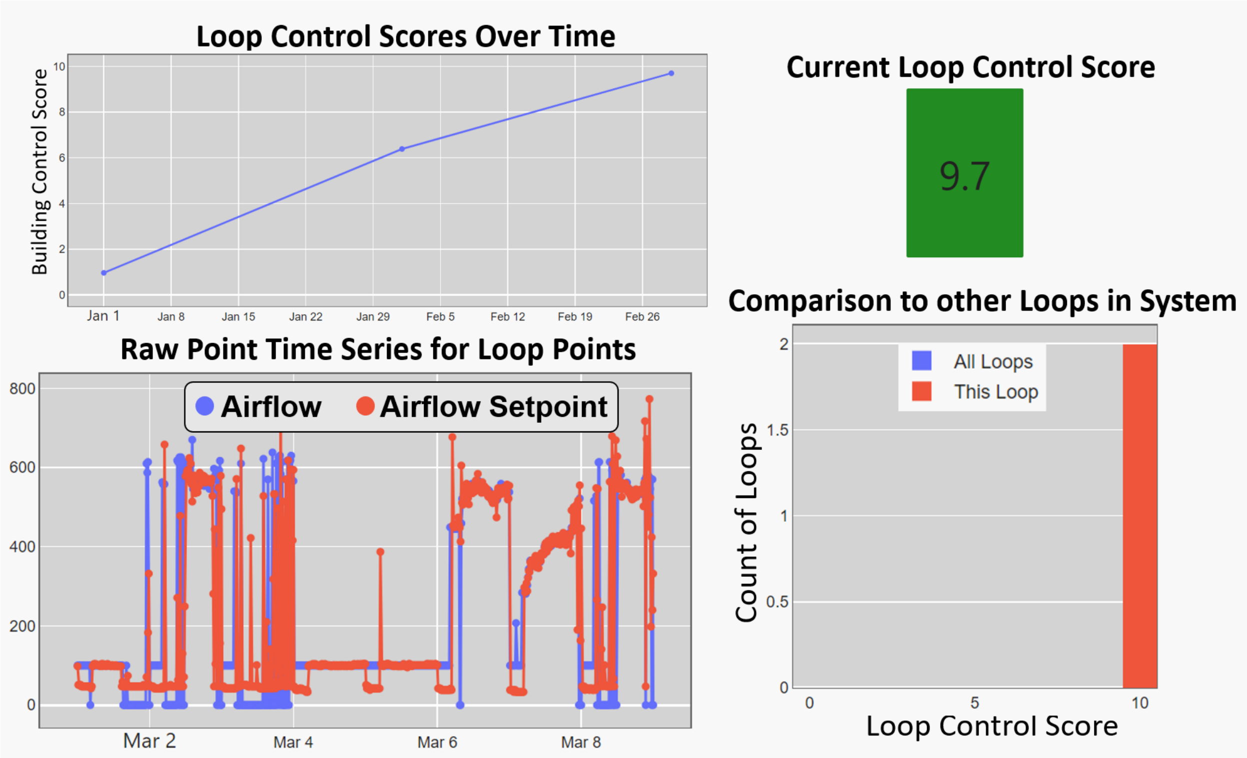

Figure 7. High Scoring Airflow Loop Performance Over Time

Figure 8 below shows the score over time and the raw data for the airflow loop in system 25. The control score over time shows a slow increase followed by a steady value around 5.5. The bottom left figure shows the raw time series data with the airflow setpoint in blue and the airflow process variable in red. The data shows that the airflow is offset from the setpoint and does not quite reach the desired values but does follow the profile of the setpoint, leading to a bad proximity and a moderate score of 4.62.

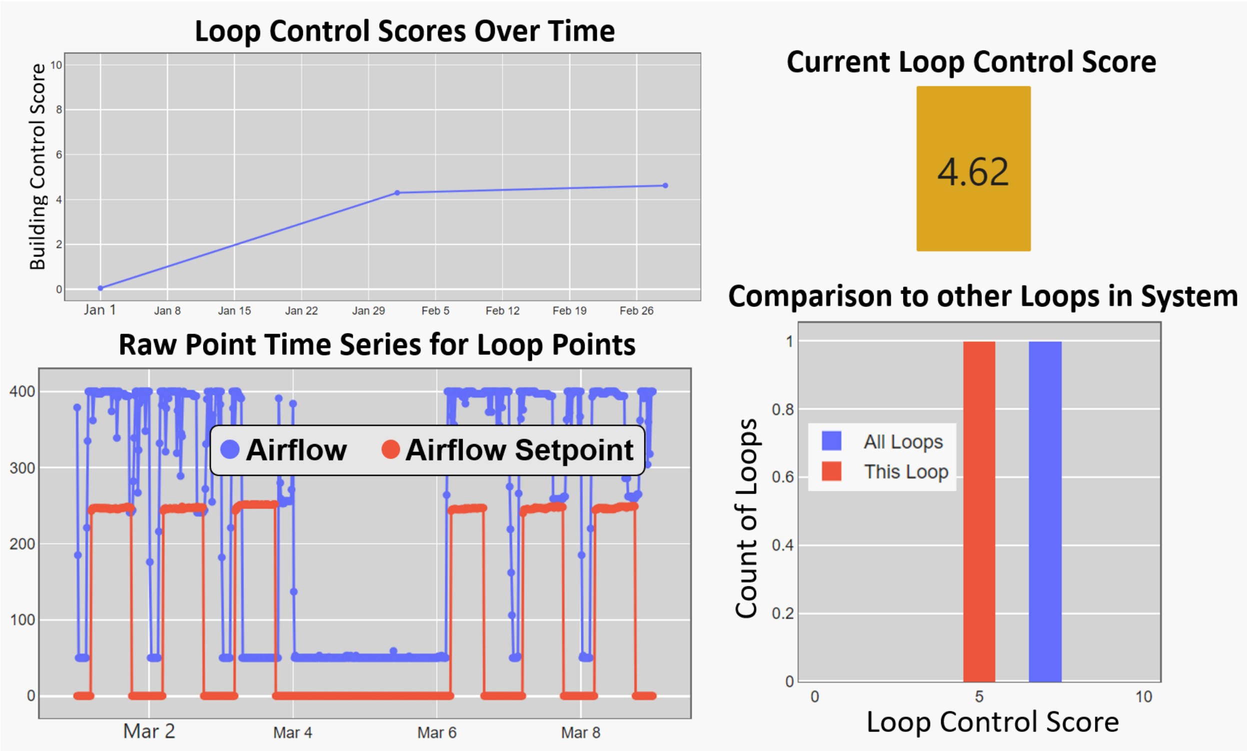

Figure 8. Medium Scoring Airflow Loop Performance Over Time

Figure 9 below shows the score over time and the raw data for the airflow loop in System 5. The control score over time shows a downward trend. The bottom left figure shows the raw time series data with the airflow setpoint in blue and the airflow process variable in red. The data shows that the airflow is 0 regardless of the setpoint, leading to a very low score of 1.

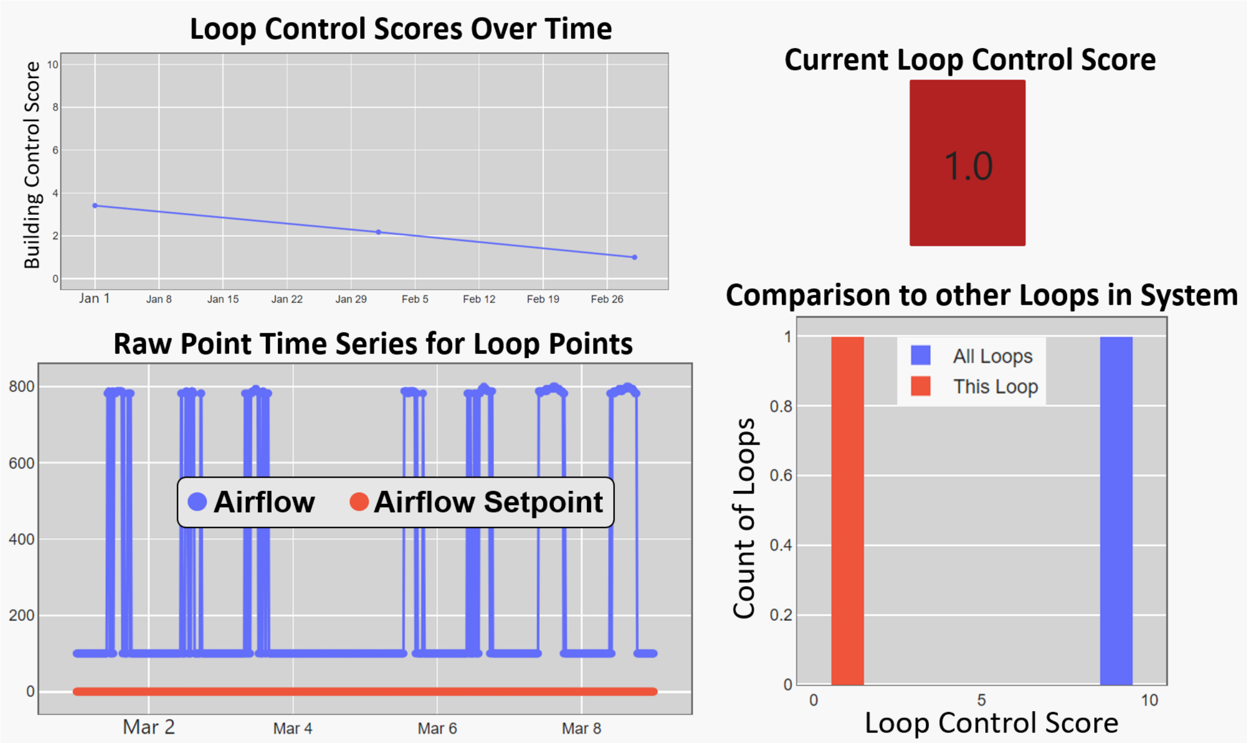

Figure 9. Low Scoring Airflow Loop Performance Over Time

2.4 Zone Temperature Loops

Figure 10 below shows the score over time and the raw data for the zone air temperature loop in System 3. The control score over time shows a steady high value around 9. The bottom left figure shows the raw time series data with the temperature setpoint in blue and the temperature process variable in red. The data shows that the zone temperature follows the changing setpoint without much overshoot or oscillation, leading to a very high score of 9.02.

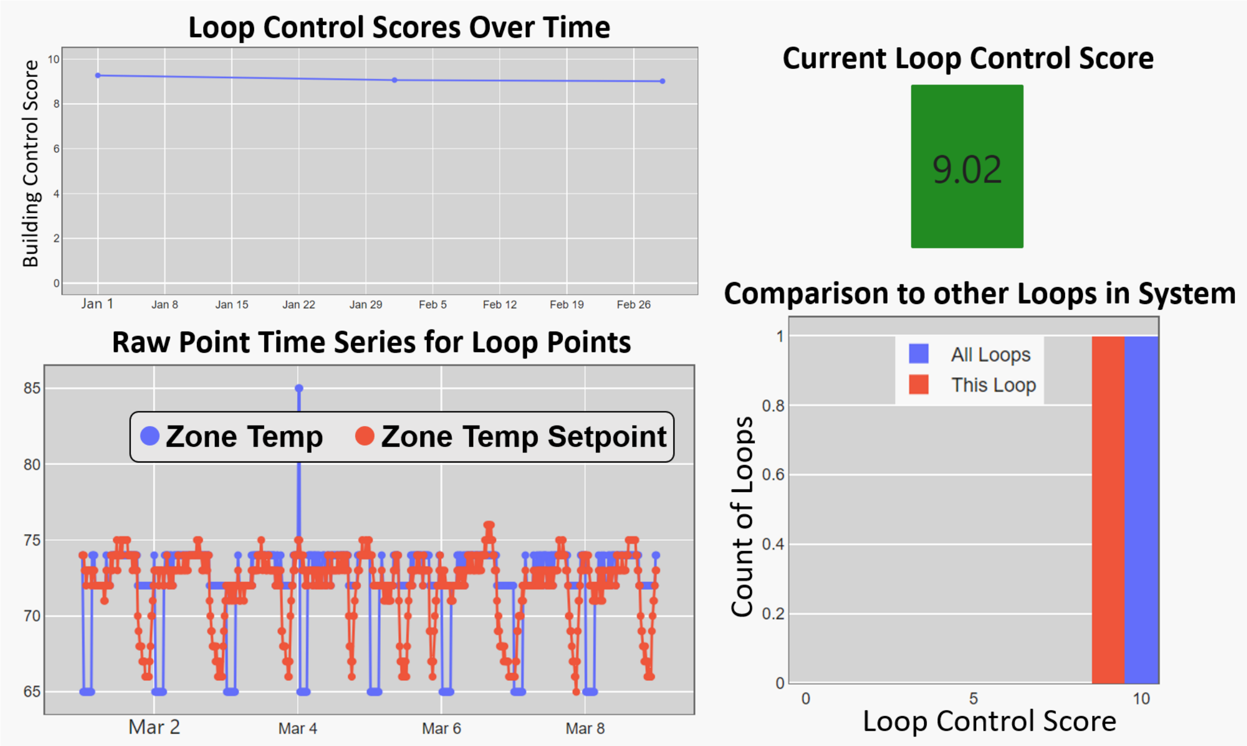

Figure 10. High Scoring Zone Temperature Loop Performance Over Time

Figure 11 below shows the score over time and the raw data for the zone air temperature loop in System 5. The control score over time shows a decrease to a steady value around 5. The bottom left figure shows the raw time series data with the temperature setpoint in blue and the temperature process variable in red. The data shows that the zone temperature follows the profile of the setpoint but often does not reach the values, leading to a below-average score of 3.83.

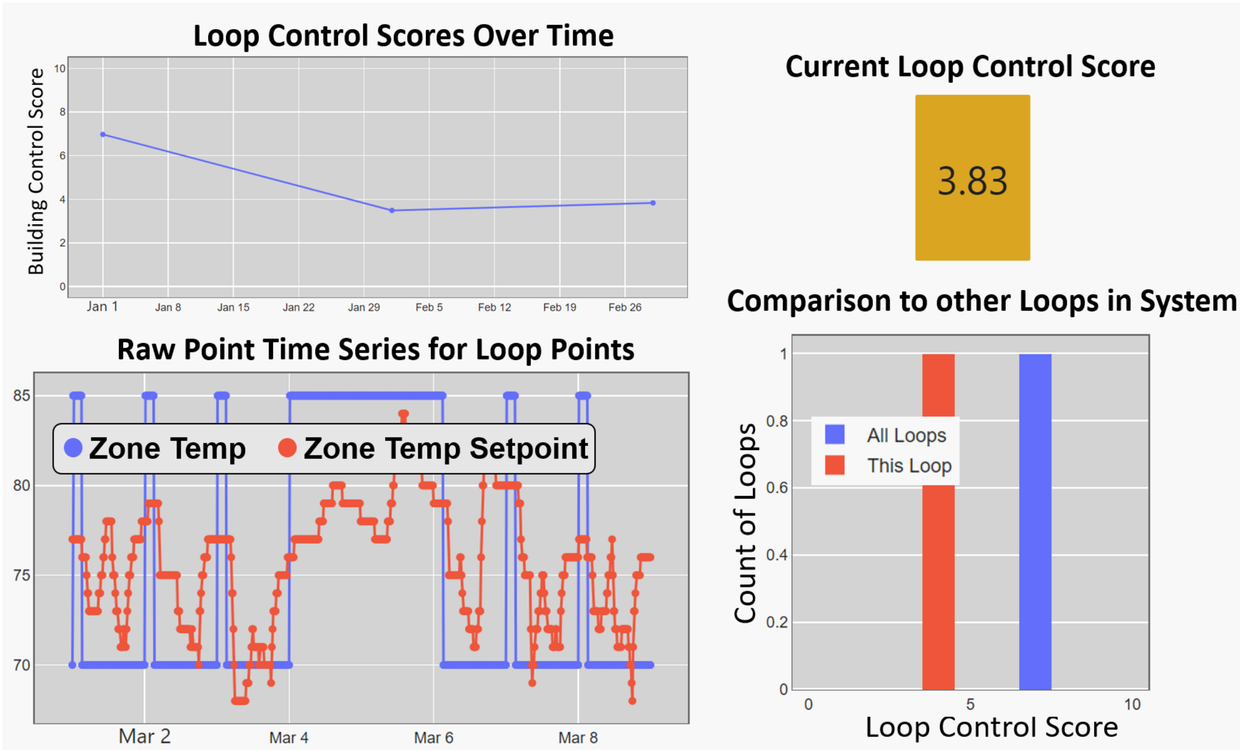

Figure 11. Medium Scoring Zone Temperature Loop Performance Over Time

Figure 12 below shows the score over time and the raw data for the zone air temperature loop in system 19. The control score over time shows a sharp downward trend. The bottom left figure shows the raw time series data with the temperature setpoint in blue and the temperature process variable in red. The data shows that the zone temperature does not track the setpoint very well and does not ever reach the desired values, leading to a very low score of 0.02.

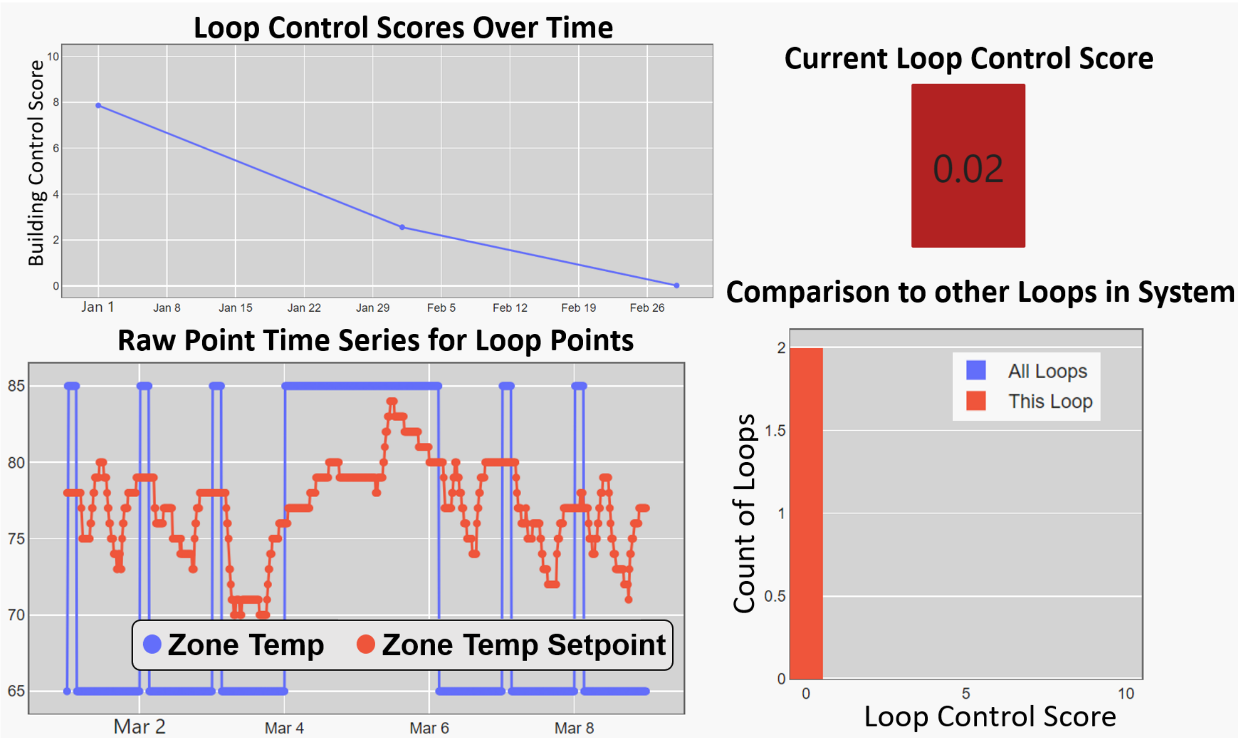

Figure 12. Low Scoring Zone Temperature Loop Performance Over Time

3.0 Conclusion

The Building ControlScore application analyzed historical control system data for the VAV systems at a GSA office building. The investigation revealed that, overall, the VAV systems exhibited a normal distribution with good control performance on average. The distribution of performance means that poor performance was observed in some systems, which should be investigated to determine the cause.

Section 3.1 provides recommendations for which systems should be prioritized for further investigation. The systems highlighted were consistently unable to reach the setpoints provided by the control systems. Investing in these issues could uncover opportunities for energy savings, improved system performance, and increased occupant comfort.

The work documented in this report has also highlighted potential areas for improvement of the ControlScore tool. Expansion of the tool's capabilities would increase its reach by reducing the time required for integration and providing the latest analysis of the buildings' control performance.

3.1 GSA Buidling Systems Identified by ControlScore for Further Investigation

Since ControlScore does not identify the underlying issues behind the poor control performance, further investigation will be required in order to fix these issues. Table 6 below shows all the loops in all systems in the building with scores less than 1. Each system has an airflow and temperature loop, but most of the low scores come from airflow loops. Some common issues with PV include positive offset (where PV is higher than SP), negative offset (where PV is lower than SP), and unstable behavior (where PV oscillates a lot). Flow loops suffering from positive offset have higher flow rates than the setpoint. Resolving this issue could lead to energy savings from reduced fan work. Temperature loops suffering from positive offset have higher temperatures than the setpoint. These issues could result from excess heating or insufficient cooling capacity. Flow loops with negative offset issues have a flowrate lower than the setpoint. These systems are not able to supply or exhaust the intended amount of air with the current equipment. Temperature loops with a negative offset have too low of a temperature and could arise from excess cooling or insufficient heating. Investigating these poorly performing loops is recommended.

Table 6. All Loops with a ControlScore <1

System ID	System Control Score	Loop	Loop ID	Loop Control Score
41	5.03	Airflow	81	0.3
4	5.02	Zone Temperature	8	0.37
51	4.84	Airflow	101	0.0
11	4.8	Airflow	21	0.69
6	4.73	Airflow	11	0.0
14	4.71	Airflow	27	0.09
70	4.7	Airflow	139	0.43
7	4.59	Airflow	13	0.01
54	4.56	Airflow	107	0.0
60	4.55	Airflow	119	0.0

System ID	System Control Score	Loop	Loop ID	Loop Control Score
50	4.53	Airflow	99	0.0
43	4.49	Airflow	85	0.0
2	4.46	Airflow	3	0.69
49	4.16	Airflow	97	0.0
63	4.15	Airflow	125	0.37
9	4.14	Airflow	17	0.95
62	4.05	Airflow	123	0.0
73	4.01	Airflow	145	0.7
56	3.91	Airflow	111	0.0
10	3.84	Airflow	19	0.48
44	3.8	Airflow	87	0.0
55	3.61	Airflow	109	0.0
27	3.36	Zone Temperature	54	0.06
45	3.24	Airflow	89	0.0
16	3.22	Airflow	31	0.0
58	2.85	Airflow	115	0.0
71	2.66	Airflow	141	0.2
8	2.55	Airflow	15	0.0
59	2.39	Airflow	117	0.0
17	2.09	Airflow	33	0.03
57	2.02	Airflow	113	0.0
15	1.53	Airflow	29	0.0
12	1.41	Airflow	23	0.27
40	1.3	Zone Temperature	80	0.74
24	1.15	Airflow	47	0.75
1	1.14	Airflow	1	0.16
52	1.1	Airflow	103	0.0
13	0.51	Airflow	25	0.0
19	0.01	Airflow	37	0.0
34	0.0	Airflow	67	0.0
34	0.0	Zone Temperature	68	0.0

ID = identification; VAV = variable air volume.

3.2 Future Work

Most of the work to deploy the tool at the new building was extracting the required metadata for the tool, as described in Section 2.1. Because the effort primarily relied on the point naming convention implemented in the historian database, it was limited and error prone. Formal metadata models, often referred to as semantic models, fully describe what each control point represents and how the points relate to each other. There are several emerging standards for these semantic models (e.g., Project Haystack, BrickSchema, and ASHRAE Standard 223P), but they have not been widely adopted throughout the building industry. ControlScore serves as one valuable example of a use case that these semantic models enable. Future work could investigate incorporating one or more of these semantic models into ControlScore and replacing the manual metadata extraction with a fully automated process. Automated metadata generation would improve the scalability of the tool by several orders of magnitude.

SkySpark, a leading energy management information system, is actively investigating incorporating the ControlScore methodology into their environment through a trial license of the tool. GSA already has a significant number of buildings integrated into SkySpark. Because SkySpark integration requires Haystack semantic model creation for each building, buildings already incorporated into SkySpark could have access to their ControlScore with very little setup required. Additional funding could be used to improve the ControlScore tool's incorporation into the SkySpark environment. With full integration into SkySpark, ControlScore could be deployed widely at GSA buildings.

Current users of the tool have provided valuable feedback on refinements that would make the tool more accessible and intuitive for users, particularly those who are not already familiar with the tool. These improvements include more explanatory material (similar to an "About" page), and documentation on how to use the tool and what the scores represent. While some of these improvements are actively being worked on, additional funding would assure sufficient staff availability for those efforts.

Finally, this deployment of ControlScore at a GSA office building is based on a static copy of the historian database for a single portion of the building. The process of exporting the static copy was done manually and would need to be completely repeated to update the scores with data from the time since the copy of the database was originally made. Such manual efforts are time consuming, error prone, and not considered best practice for similar analytic tools. Future funding could go toward provisioning the required cloud infrastructure for an automated data pipeline that would regularly pull in updated data for the tool so that it could provide up-to-date insights on the building's ControlScore.

4.0 References

Salsbury, T. I., A. P. Rogers, T. A. Yoder, S. R. Johnson, and X. Duan. 2023. “A New ControlScore Concept for Building Performance Assessment.” *Journal of Building Engineering* 66 (2023): 105770. https://doi.org/10.1016/j.jobe.2022.105770

Pacific Northwest National Laboratory

902 Battelle Boulevard
P.O. Box 999
Richland, WA 99354

1-888-375-PNNL (7665)

www.pnnl.gov