

Dashboard for z/VM Energy Consumption

z/VM Performance Data Pump

Behind the Dashboards

April 2025

Rob van der Heij
robvdheij@nl.ibm.com

z/VM Development – Endicott, NY

z/VM Energy Consumption Dashboard

Power Usage metrics introduced with z17

- Requires suitable level of z/VM to gather the data
- Updates to z/VM Performance Data Pump

Sample Dashboard

Sample dashboard will be provided to visualize:

- Power usage of the entire CPC
- Power usage apportioned to z/VM
- Breakdown by workload class and by Linux guest
- Power usage over time and aggregated over 24 hours

Knows Restrictions

- Apportionment to Linux guests is an approximation
- In a shared resource environment, apportionment of resource usage to individual resource consumers is an approximation since workloads will affect each other

Disclaimer

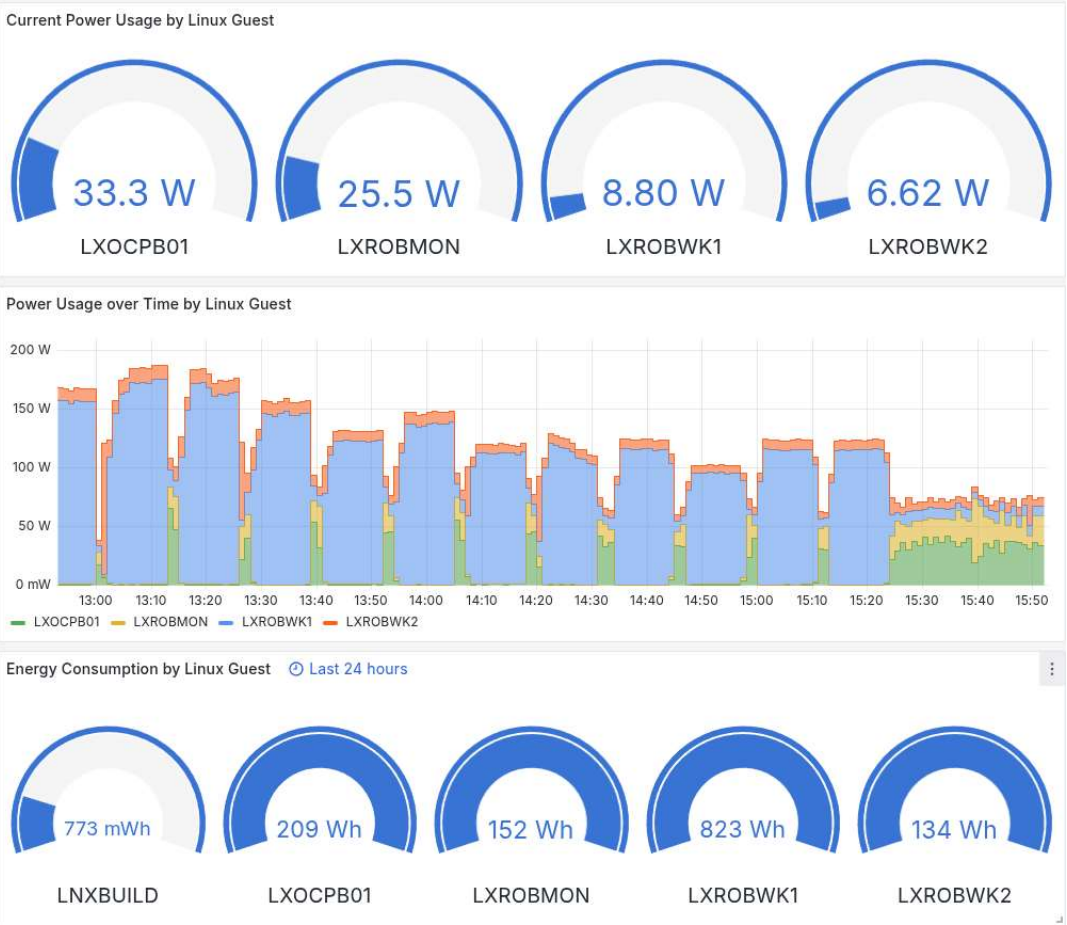
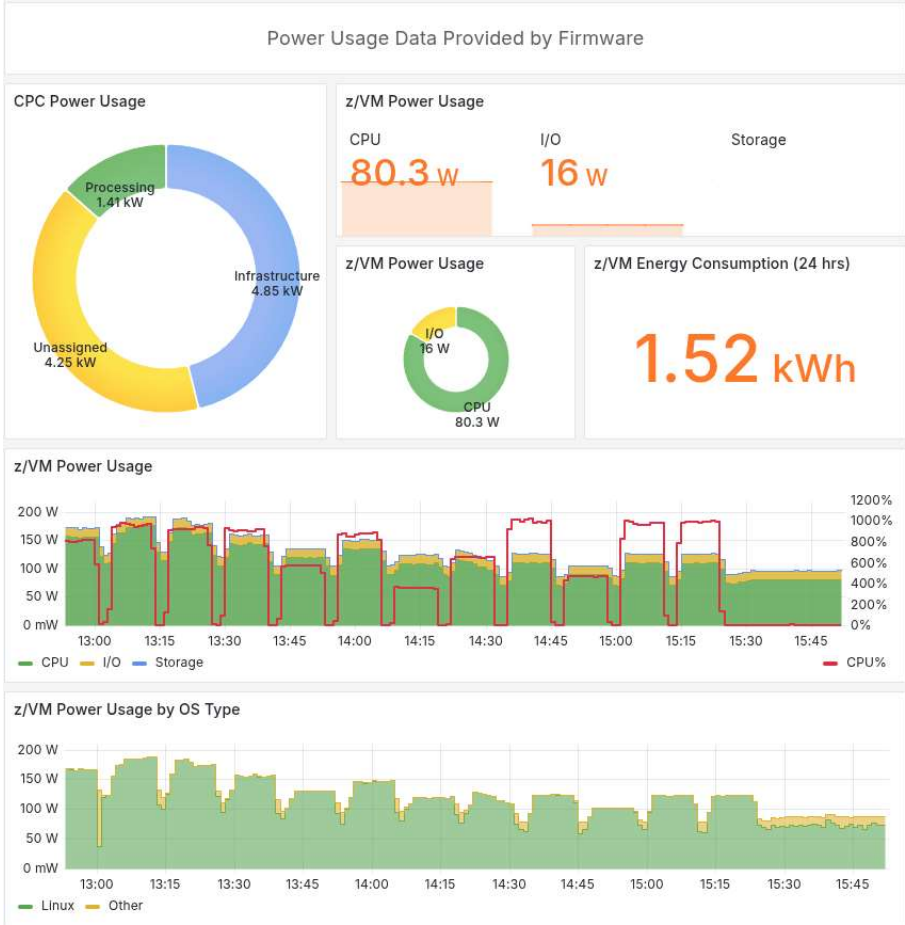
- This is a sample dashboard; clients must verify whether it is effective for their configuration and workload

This dashboard is licensed by IBM under the Apache 2.0 License and is provided 'as is' without warranty, representation, support, maintenance or an obligation to issue updates.

z/VM Performance Data Pump

- z/VM includes the power usage metrics in the z/VM monitor data stream when provided by the firmware
- The z/VM Performance Data Pump is included in z/VM Performance Toolkit product that is pre-installed on z/VM but separately licensed
- For z/VM 7.3 the Performance Data Pump is delivered with VM66765 UM90432
- For z/VM 7.4 the Performance Data Pump is included in the base release
- Support for the power usage metrics requires additional service for Performance Data Pump
- To take advantage of the sample dashboard, the client must provide an InfluxDB and Grafana service (open source projects).
- The "Grafana Dashboards - Fast Start" publication has detailed instructions for installation and configuration of InfluxDB and Grafana (whether on s390x or on x64) <https://www.vm.ibm.com/related/perfkit/datapump/>
- The IBM Z Container Registry provides s390x images of both InfluxDB and Grafana for clients who prefer to deploy the services in a Linux guest on z/VM
- PERFKIT screens and reports will not be enhanced to show the power usage metrics

Dashboard Overview



Dashboard Navigation Hints

The [system] menu can be used to select the z/VM system to show in the dashboard. The menu requires no configuration and will automatically list the systems that provide the relevant data.

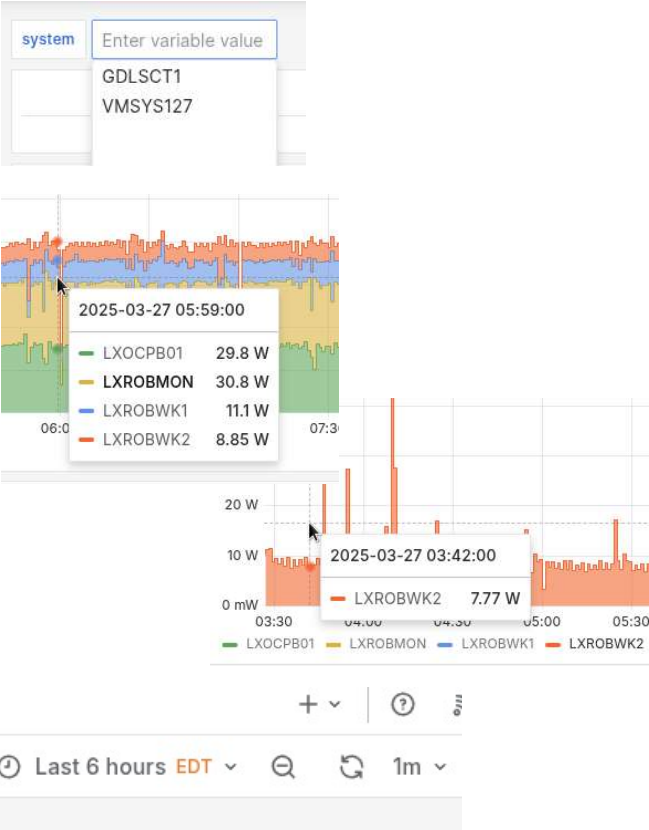
Hovering over a chart with the mouse normally displays the numerical values from the chart.

The menu in the upper right corner of each panel can be used to view a panel in full-window mode to study details in the data.

For mixed charts, selecting the label in the legend shows a series in isolation or excludes the series from the chart. This is useful to a dashboard where one series dominates the display.

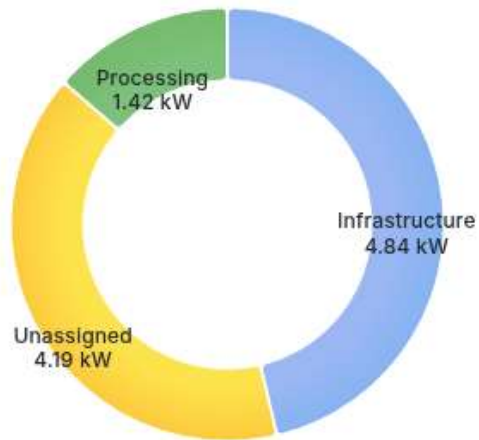
The default dashboard shows the most recent hours of data. The time range can be selected through the pull-down menu in the upper right corner. Select a part of the chart to zoom in on a shorter time range.

The dashboard can be saved in Grafana to keep a different time range or selection. The Grafana profile can be used to select either light or dark style.



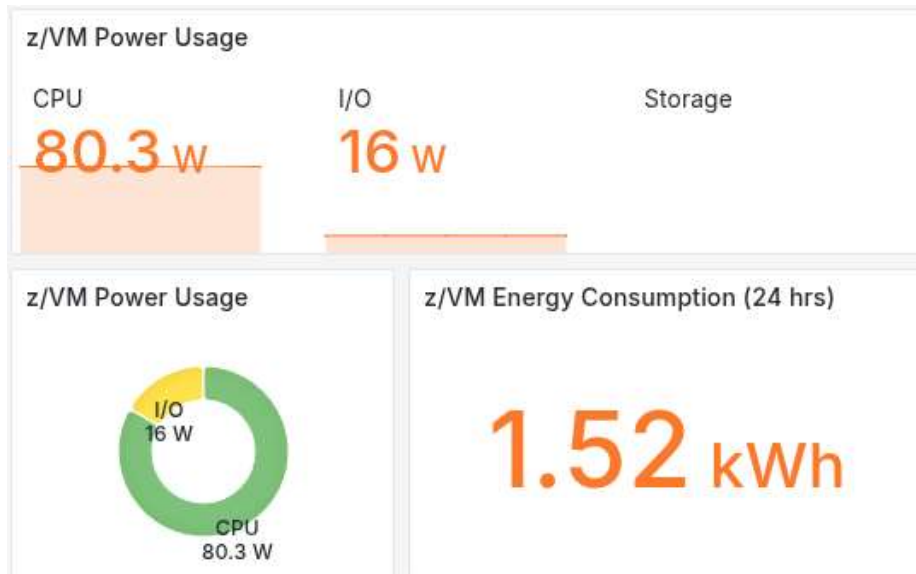
CPC Power Usage

CPC Power Usage



- The chart shows the breakdown of current power usage of the entire CPC where the z/VM LPAR is running.
- Like any electronic device, a portion of the power usage of an IBM Z configuration is fixed and does not depend on the workload.
- For a shared resource environment like IBM Z, efficiency is in sharing the resources among the workloads rather than suspending hardware components when the workload reduces.
- **Processing** - The electronic components like CPU, memory and I/O
- **Infrastructure** - Additional components beyond processing
- **Unassigned** - Any parts in standby or reserved state
- Because these charts are from an internal lab environment, the portion unassigned capacity is likely much higher than for a typical client situation in which the workload is consistent and excess hardware capacity is avoided for economical reasons.
- When different z/VM systems on the same CPC report the CPC power usage, the numbers can vary slightly because of the timing of the z/VM monitor.

z/VM Power Usage



- These panels show the power usage apportionment to the z/VM LPAR, based on the LPAR configuration and resource usage.
- The top panel shows the breakdown of current power usage by CPU, I/O, and Storage (memory). The shaded are show an approximate trend.
- The pie chart shows the same breakdown, but in a visually simpler manner.
- The panel at the bottom right shows the aggregated usage of the z/VM LPAR over the past 24 hours.
- The aggregated usage is based on hourly averages while z/VM was running.

z/VM Power Usage over Time



- Two charts show the apportioned power usage for the z/VM LPAR over time.
- The top chart shows the breakdown by resource (CPU, I/O, and storage). It also shows a line chart with the overall CPU usage of the workload (where 100% is one full CPU of capacity)
- The bottom chart shows the breakdown of power usage by type of workload, based on CPU consumption. For many configurations we expect the Linux guest workload to be the dominant factor, but guests running z/OS or z/VM would also show here.
- The power usage of the z/VM LPAR is apportioned to the workload type based on the user CPU usage. This means that any z/VM system overhead is also proportionally charged to the virtual machines.

Power Usage by Linux Guest



- These two charts show apportionment of the z/VM power usage by Linux guest.
- The top chart shows the current power usage by guest, with gauges allowing easy comparison between guests. The bottom chart shows the power usage over time as stacked graph, revealing both total usage and breakdown.
- The breakdown by guest is less realistic for an almost idle z/VM system because the fixed portion of the power usage is charged to the individual guests based on almost no usage of CPU resources.
- Stopping or moving a Linux guest may not reduce the z/VM power usage by that same amount because of shared resources and apportionment of idle power consumption.
- When the number of guests is too large to for an effective presentation by guest, it is advised to group the Linux guests by some attribute like account code. This is possible with small changes in the sample dashboard.

Energy Consumption by Linux Guest

Energy Consumption by Linux Guest 🕒 Last 24 hours



- The panel shows the aggregated energy consumption by Linux guest over the past 24 hours.
- The energy consumption is based on the usage of CPU resources by the guest during the past 24 hours.
- In a shared resource environment, clients should not expect that moving a Linux guest to another LPAR would be apportioned the same amount of energy consumption.

IBM