

Optimising Capacity and Performance Management for Cloud Services

Benjamin Wiriyapong

MSc in Big Data

BrightSolid

BrightSolid is a leading UK cloud services provider. The company offers many cloud service solutions. Recently, the company has invested in Hyper-V platforms which allow customers to fully customise their cloud resources. However, as Hyper-V is a multi-tenanted cloud platform, it is important for the company to know how the resources are utilised and when there is a need to add more resources.

Objective: This project aims to find metrics to indicate the capacity of Hyper-V platform, particularly memory utilisation.

Data Source: System Center Operation Manager (SCOM)

System Centre Operation Manager is an infrastructure monitoring tool from Microsoft. The data monitored is stored in:

- **OperationsManager:** This is a SQL server database that contains recent configuration data.
- **OperationsManagerDW:** This is a SQL Server database that stores monitoring data for historical purpose.

Methodology

In order to come up with metrics and relevant data, a meeting with internal stakeholders was held to define their needs. Then data sources were identified.

It is recommended by Microsoft to work with views because the data is more abundant than tables. There are 3868 views in 'OperationsManager' and with experience, their names are rather self-explanatory. 'OperationsManagerDW', on the other hand, has 176 views but it is quite difficult to work with as relations between them were unknown. With assistance from engineers that work directly and more experienced with SCOM, documents from Microsoft on the System Centre site; and an ability to execute SQL, relevant views were found e.g. Perf.vPerfRaw, vManagedEntity, vPerformanceRuleInstance, vPerformanceRule.

Key Performance Indicators

As the purpose is to identify metrics and threshold that can be used to identify a host's capacity, there are five areas that should be looked at in order to indicate the health of the hosts:

- Hard Disk
- Memory
- Processor
- Network
- Process

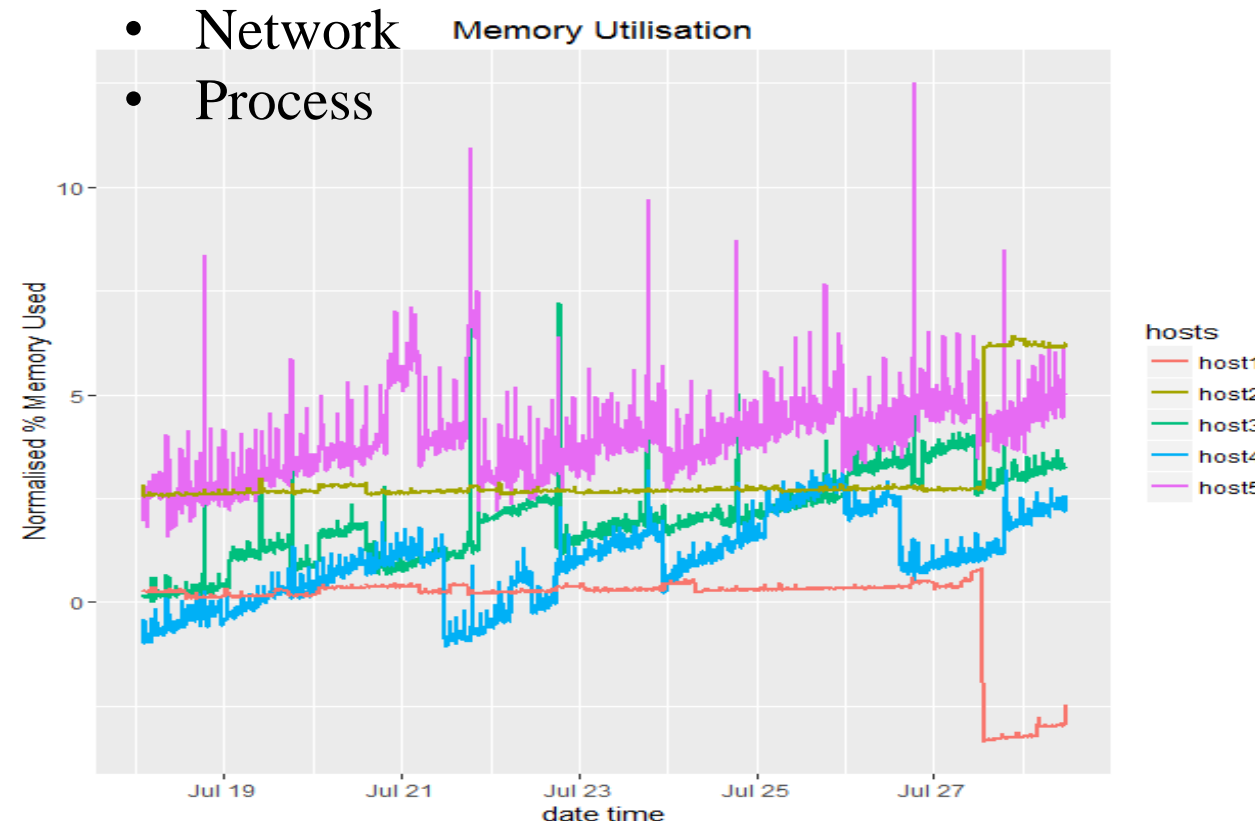


fig.1 shows normalised % memory utilisation



Out of those five areas, the primary concern of the company is the performance of hosts. Therefore, the area that shall be paid particular attention on is memory as it indicates the performance capacity of the hosts. The following metrics in SCOM which measures memory are:

- Committed Byte in Use
- Available Mbyte
- Free System Paged Table Entries
- Pool Non-Paged Bytes
- Pool Paged Bytes
- Pages Per Second

Management Pack

Management packs define the information the SCOM can collect, evaluate events and operations to ensure health of the servers.

The counters can be captured if the management pack is installed. The current hardware that Hyper-V is running on can only provide the 'Available Mbyte' metric. That seems to be the most suitable counter for this project. The visualisation (fig.1) shows the memory utilisation on each host. From experience we know that the decreases are caused by VMs migration between hosts to maintain the health of the hosts.

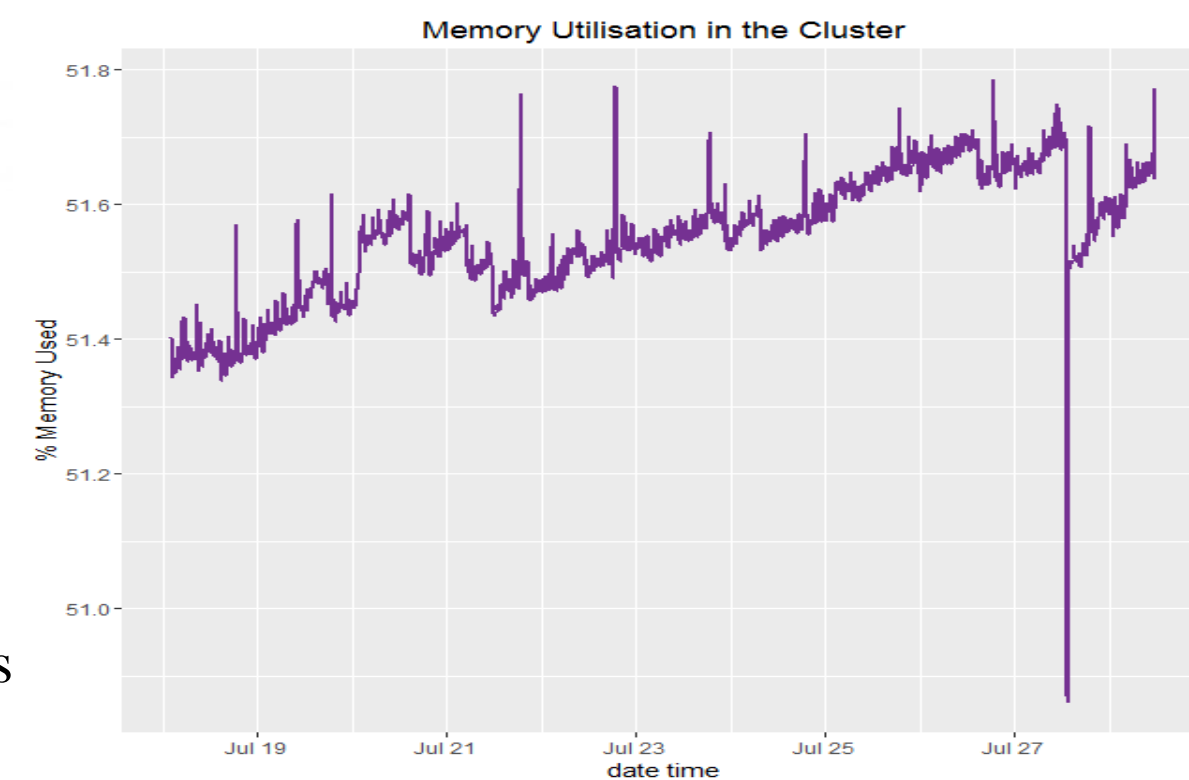


fig.2 shows % memory utilisation of Hyper-V cluster

However, the aggregation of the memory utilisation does not exceed the 80% threshold (fig.2). The status of the hosts are healthy as this can be verified from 'OperationsManager' database and SCOM console.

Next Steps:

Model building is attempting to forecast when the host will reach certain threshold and will eventually impact the performance. Normal equation linear regression and time series are being considered. The best fit model will be integrated to a dashboard with relevant metrics to assist engineers to monitor capacity. Further as the environment is dynamic, more VMs can be deployed on daily basis, an approach to develop the model to online-learning is also taken into account