

Maximizing the use of excess capacity on your IBM System System i™ or IBM System p™

PM for System i or PM for System p and the IBM Systems Workload Estimator

Introduction:

The effective use of IT systems in today's business world has never been more critical. The 24x7 global business environment demands it. Senior company management demands it. It is one area of business operations where companies must not only keep pace to remain competitive, but is also an area that can allow them to 'get ahead'.

So what are the implications of this changing world on the IT department?

The role of the IT manager is changing. Maximizing return on the investment in IT is the 'name of the game'. They need to be more innovative. The need to bring technology to the forefront has never been more pressing...as it drives business results. The day to day challenge to run the business more effectively while also responding to the growth of the business is very real. IT executives must have a view of IT operations so that they can manage the growth and risk of operations, while at the same time, understand the system implications of adding 'game changing' strategies for using their IT investment.


It is sometimes easy to imagine what else you could be doing with your system to really help the business but it can be really challenging to try and assess if your current system has the capacity to implement the new strategies. Or if instincts indicate a capacity shortage would exist, it is difficult to assess what additional investment would have to be made to accommodate the new direction.

So what can help?

IBM's Performance Management for System i (PM for System i) and Performance Management for System p (PM for System p) in conjunction with the IBM Systems Workload Estimator (WLE) can give you a head start on helping understand the possibilities of what you can do with your 'excess' capacity to help address the new opportunity requirements. These tools are readily available to nearly everyone wishing to take advantage of them. Terms and conditions apply, (see below) but most customers can use these tools for no additional charge.

PM for System i and PM for System p (PM) are normally thought of as 'headlights' into helping clients understand approaching capacity issues. They are in fact excellent at that. The trend analysis that is calculated and graphed against remaining available CPU and DASD capacity give a good first view of any impending capacity constraints. That helps make it easier to plan for needed system and associated budget requirements.

Peak Processor Utilization, Total Three month trend, the two heaviest hours per measured day



The CPU trend graphs provides a view of utilization trends over the past 13 months with a projection of utilization growth based on trends for the last 3,6, or 12 months.

Identifying excess capacity

However PM is just as useful in pointing out to users the amount of ‘excess’ capacity available to them on the partition or systems monitored. You can easily see by hour, day, or month how busy your system really is in comparison to its available capacity. And with the special capability to integrate the collected historical PM data with the WLE, you have the information to help begin assessing what other kinds of things could be done with the system in addition to being able to assess what changes (e.g. upgrades, extra processor) would need to be made if the system was in a capacity constrained status.

With excess capacity identified, the PM to WLE process can help you evaluate potential new uses of your system in a number of areas. For example:

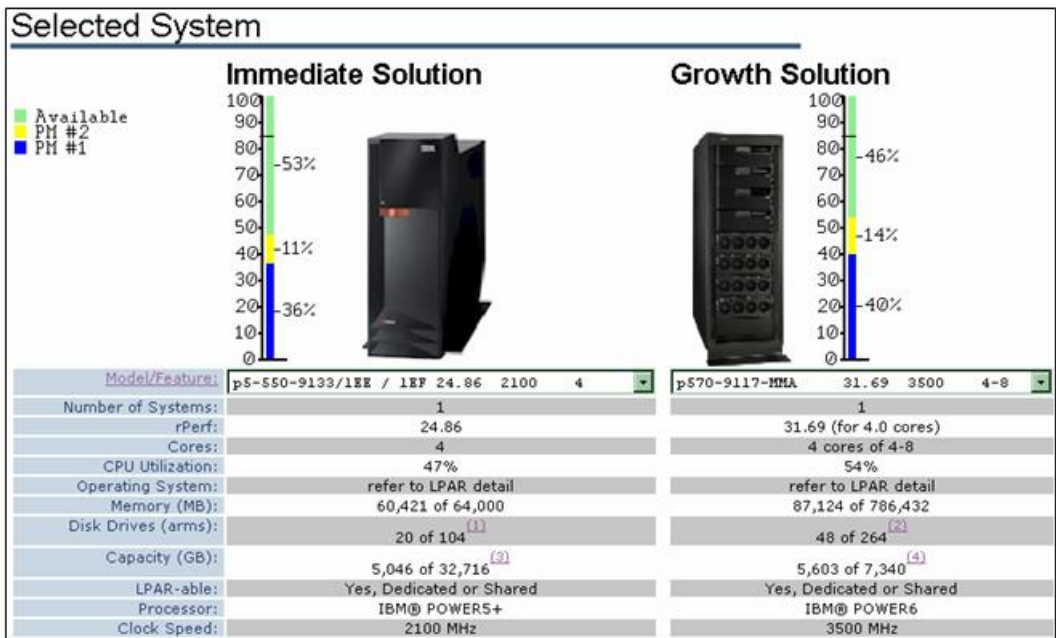
- a server consolidation...bringing workload over to the System i or System p from another system to take advantage of the excess capacity
- the addition of another partition(s) to run a special new application(s) ...whether it is an i5/OS®, Linux® or AIX® application
- the implementation of additional software technologies, like IBM middleware (e.g. Domino®, WebSphere®, IP Telephony)

The WLE recommendation

The WLE helps by providing a high level projected view of what your combined processor utilization will be while supporting the new desired workload plus the current system’s workload (as reported by PM). As a default, the WLE will automatically recommend the next larger upgrade or replacement system available in IBM’s lineup of systems that will handle the combined work. However, the user has control over the selection of the system that the combined workloads are projected on. For example and as applicable to the scenario where the PM data is indicating ‘excess’ capacity, the user can easily select from a dropdown of IBM systems the machine type and model that most

closely approximates their current system configuration. WLE will then ‘rerun’ the sizing against the designated system.

The WLE also automatically indicates the system (CPU, memory, DASD, and DASD arms) that is needed right now to support the combined workloads along with a system that will be needed a year from now. The ‘one year out’ view is known as the growth recommendation. The one year projection is based on the growth rate that the historical PM data has indicated for the growth of the current workload plus any indicated growth for the new application. At user option, the growth option can be set to project the system requirements for up to 3 years. This can be very valuable when planning as it is often useful to know if and when future upgrades may be required.



The Workload Estimator considers the utilization and growth trend information provided by PM for System p (or PM for System i) when recommending the upgrade needed immediately (left column) versus one year from now (right column).

More detail on how PM – WLE sizing works

When clients are viewing their PM graphs on the Internet, either the ‘no charge’ summary level graphs or the nominal charge detailed level graphs, there is an option available to ‘Size next upgrade’. There is no additional charge for use of the PM – WLE sizing function. Clients may use it as often as they like. It is highly recommended that

users take a few minutes to review the online Help tutorials in the WLE to understand its operation.

When the 'size next upgrade' option is selected, customer historical PM data that is securely stored at IBM for the selected system is automatically loaded over to the WLE into the Workload Definition screen. There is no need to key anything to get the data loaded. From there, if one takes all the 'default' options, the WLE will simply show the customer's current workload projected on the next System i or p that is either in the upgrade path or is large enough to handle the projected workload. Users also have the option on the Workload Definition screen to manually manipulate growth rates, etc. to alter the projected outcome.

To add another workload to the configuration, simply select the Workload Selection tab. WLE then presents a window where you can 'add a partition' (and select the OS type to run in the partition) and then 'add a workload'. Under the Add a Workload tab, a dropdown of applications to choose from is presented. Choices include everything from adding a 'generic' workload (where you are subsequently presented a 'generic' template) to IBM middleware templates (e.g. Domino, Workplace™, WebSphere). The templates present key questions for user response that help describe the characteristics of the new workload. The WLE will then consider the PM workload plus all newly described workloads in the projected sizing.

Should the WLE indicate that the selected system is not large enough to handle the desired workloads; it will suggest a larger system. And if the system is a Capacity on Demand configuration and has available processors that could be assigned to the system or partition, the user can select the 'Modify / Visualize system' option. It will then provide an opportunity to add another whole or partial processor to the system or partition and then reveal the change in utilization percentages. Therefore it is easy at any time to understand the impact of adding a Capacity on Demand processor.

Capability also exists to combine the PM workloads from multiple systems or partitions all into one existing or new system. Simply indicate in the PM screens which systems to combine and WLE does all the rest of the work for you.

WLE recommendations can then be printed, saved to PDF for easy transfer or saved to your hard drive for later restoration in subsequent sizings.

It is important to remember that the recommendations made when using the PM to WLE tools process are an estimate of sizing needs. IBM recommends that you always review the WLE recommendations with either your IBM representative or IBM Business Partner before making any final purchase decisions.

So why PM for System i and PM for System p?

The key to all this capability is having a continuous record of your system's current utilization and growth. That is the absolute basis needed for sizing any future needs.

This applies whether the system environment is capacity constrained or if the system has excess capacity and there is interest in being able to understand what other value can be received from the system. PM helps minimize a lot of the guess work on your existing environment and helps ensure that an accurate reflection of your current utilization and growth trends are being used to assess future sizing options.

Terms and Conditions

PM for System i and PM for System p both have ‘no additional charge’ and ‘nominal charge’ reporting options. The no additional charge option provides a summary level graph that is updated monthly and accessible on the Internet. The graph includes utilization measurements on key metrics on the CPU and disk...and projects the growth rate for the coming twelve months against the capacity of the system or individual partition being monitored. Integration of the PM data with the WLE is also available for use as often as desired. To be eligible for the no additional charge option, the system must either be on warranty or be covered by an IBM hardware maintenance agreement.

The detail graph reporting option provides up to 30 different reports and graphs, many of which are updated daily. Clients have access to an ‘interactive graphing function’ option that allows the user access to up to 24 months of their stored historical PM data at IBM Rochester, Minnesota, for creating selected graphs over the time period of their interest. This can be valuable in multiple ways. As example it enables an ability to do a limited amount of problem diagnostic drill down to a particular time period of interest to aid in better understanding of a problem. The detailed offering is available from IBM Global Services. See the ‘contacts’ topic in the Overview tab of the PM for System i and PM for System p web sites for more information on who to contact in your country for applicable charges.

Next steps:

If you are not using PM for System i or PM for System p already, then get started.

Simply follow the ‘Getting Started’ instructions on the respective web sites:

PM for System i: <http://www.ibm.com/systems/i/pmsystemi5>

PM for System p: <https://www.ibm.com/systems/p/pm/>

A Redbook is available on the referenced web sites for explaining in more detail the graph options on the no additional charge summary report option or the nominal charge detail report option of PM. The web site and Redbook also describe for you the PM process of data collection and data transmission to IBM that enables the creation of the PM performance and capacity reports.

IBM, the IBM logo, System i, System p, i5/OS, AIX, Domino, WebSphere and Workplace are trademarks or registered trademarks of International Business Machines in the United States, other countries or both. Linux is a registered trademark of Linus Torvalds in the United States, other countries or both.