Migration to POWER9

How to make a better sizing with LPAR2RRD

Migration Drivers

Every migration is risk by nature, and usually is driven by something. In these days is most common driver is money. Every single IT Manager facing to “cost reduction” challenges which need to handle. In case, you are running IBM Power Systems or any other Unix systems, migration to POWER9 could be your solution. Reasons for migration to POWER9 may vary from just replacement of deprecated hardware to multi-platform migration.

In any case, purpose of the new server has to be carefully considered from many points of views. In example another type of server will be used for web/application server and another for large database in mission critical environment.
As there is so many thinks to be considered and appropriate setup has to be chosen, there is no universal design in place which can apply to all solutions.

Technical constrains

There is list of main points, which should be taken into account in case of migration planning:

Reliability, Availability, Serviceability - It will be possible from business point of view to switch server off for maintenance easily? It is just part of application farm and could be easily switched off or it is mission critical application in single instance? Do I need high-availability and on which level? What about DR solution and disaster recovery plans? How much cost me, if this application is down and on other hand, how much cost me, to keep this application running? Should I take live partition mobility into account?

I/O adapters - Which kind of I/O operations will be handled and what about volume? All adapters could be virtualized and shared or do I need dedicated adapters for any reason such as heavy load on SAN or LAN? Alternatively do I need isolated/dedicated adapters for security reasons (corporate rules to isolate environments/landscapes and so on)? Where I can put my boot disks - internal disks, VIOS on local disks, VIOS on SAN? Where I can put my swap and what is proper size? Can I use my multi-path drivers and what are they requirements?

CPU workload - How will target load looks like? Which benchmark I should use for proper sizing? Especially in case of migration from another platform proper analysis has to be done. Where peaks of utilization are in particular LPARs to maximize efficient load on physical machine and on other hand not overload physical machine? Which SMT should I use? There could be significant difference in performance base on SMT settings, type of application, application settings and type of workload.

RAM - How much memory I will need? In case of migration from another platform proper analysis has to done to ensure difference in memory consumption on Power platform.

Security - Should I run all defined LPARs on one physical box? Is provided security certification for Hypervisor sufficient? Do I need any extra security features?

Compatibility - In case of new physical server installation, every time has to be considered required minimal version of operating system. One thing is to have appropriate OS level for installed hardware and firmware levels, bud second thing is to have appropriate or certified version of OS for running application. Do I have adequate version of HMC software? Should I install required version on current HMC console hardware?

Memory, security, compatibility, HA and DR are all defined but what POWER9 model to choice with how many CPUs to be safe? This one, should be really difficult. Do I'll need upgrade capacity of the server for future? How big? Small entry-level and midrange servers could be extended by second CPU or up to 4 CPU depending on server type. Some of them cannot be upgraded at all. On small servers all installed resources are automatically activated. On other hand these servers are cheaper than enterprise class servers, which allow incremental HW upgrade and also providing CuoD (Capacity upgrade on Demand) for CPU cores and MB of memory.

And finally there are licenses of running software which has to be considered. Also from many points of views as is listed below:

- Licensing model - per core, per user, etc ….
- Price of license per unit
- Performance per core
- Licensing tier per server class i.e. PVU - blade and entry level, midrange and enterprise class have different PVU values
- etc

How to get guidance for sizing

In case, you are considering migration from older Power CPU to newest one – POWER9 for whatever reason, you can easily plan this migration with LPAR2RRD using feature CPU Workload estimator (CWE). Especially this feature helps you with projections of your current (and/or historical) workload which you have on current infrastructure, to a new machines.

Benefits with using LPAR2RRD CWE is really easy projection of current workload (and/or historical) to new physical box or boxes based on your real workload data already collected by the tool. In this tool you can easily select current servers considered for migration. These servers should be whole physical boxes (one or more) or only one or more selected LPARs. This allows you, to make many different scenarios for planned migration and consolidation without big effort. You can plan migrations and try to fit to as much types of servers as you can.

Another help could be with finding of proper I/O adapters sizing. If you need to know about data throughput generated by whole physical server or just some set of LPARs you can create your own Custom group and see aggregated throughput on this set of LPARs. Based on this information you can choose adequate number and types of I/O adapters. Even if you considering some kind of consolidation, you can create Custom group contains LPARs from more one physical server or even contains all LPARs from few physical servers to see appropriate result.

You can create for example these custom groups for proper I/O adapter sizing:

- Whole physical server
- Set of all production LPARs
- Set of all non-production LPARs
- Set of all DB LPARs
- Set of all LPARs across few physical servers

For these custom groups then you can see overall utilization of all resources from CPU and RAM to Ethernet and FC.

On following pictures are sample of CWE usage.
On first picture you can see LPAR/Physical machine selection:

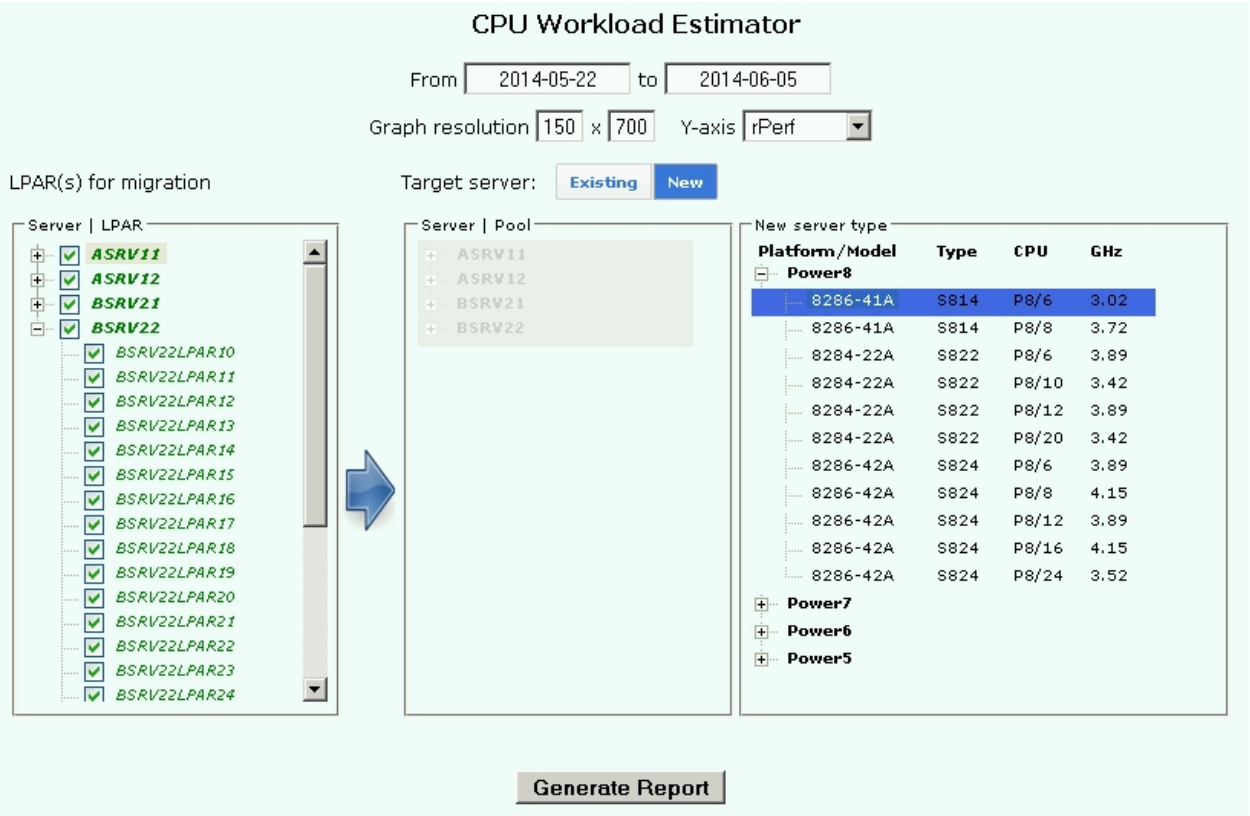

On next pictures you can see result. Result was cropped due to lenghts cause 100 of LPARs:

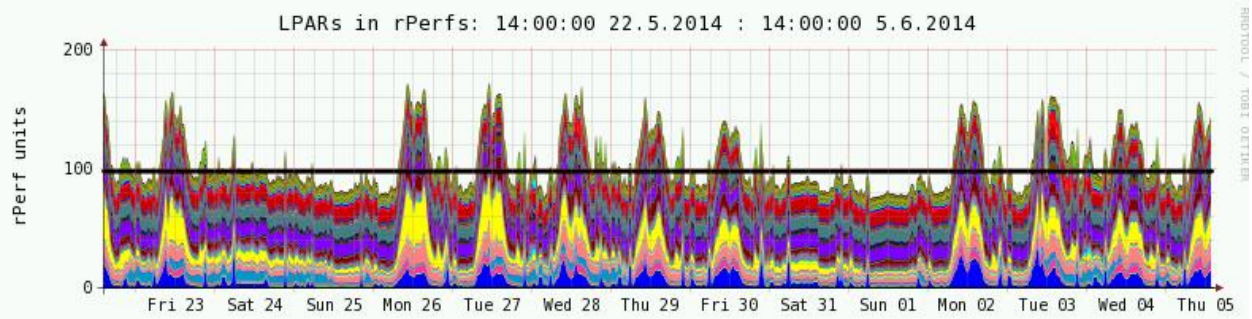

Server - LPAR (will be migrated)	average	maximum	[rPerfs]
ASRV11 - ASRV11LPAR10	5.9	35.4	
ASRV11 - ASRV11LPAR11	0.4	1.0	
ASRV11 - ASRV11LPAR12	0.4	1.0	
ASRV11 - ASRV11LPAR13	2.3	7.5	
ASRV11 - ASRV11LPAR14	0.4	1.0	

CPU limit for target server:
IBM Power S814 (model 8286-41A) 98 rPerfs

Server details	number of cores	GHz	rPerf/core
IBM Power S814 (target)	6	3.02	16.2
ASRV11	16	3.0	9.7
ASRV12	16	3.0	9.7
BSRV21	16	3.0	9.7
BSRV22	16	3.0	9.7

Conclusions

This tool could help you with sizing by calculating performance based on rPerf/CPW benchmark, nevertheless this is only one of the parameters which accuracy depending on accuracy of rPerf/CPW benchmark itself. On other hand, this tool should be used as agentless without installation on LPARs, very easy usage, reasonable outcome as it is based on YOUR real data.

Real migration result could be affected by things such as OS levels, OS setting/tuning, firmware levels, type of load (single threaded vs multi), limited by rPerf/CPW accuracy, ability of application to use new CPU and SMT features and many others.
Despite all of that, how else you can make easily POWER5/6/7/8/9 comparison?

Anytime, our team is ready to help you to apply our common sense, best practice and experience acquired on many other customers.