

Capacity planning for IBM Power Systems using LPAR2RRD

Agenda

- LPAR2RRD and STOR2RRD basic introduction
- Capacity Planning – practical view
- CPU Capacity Planning
- LPAR2RRD
 - Premium features
 - Future
- STOR2RRD quick introduction

Basics of LPAR2RRD and STOR2RRD

- Free performance monitoring tool for IBM Power Systems and storages
- It is an OpenSource distributed under GPL v3
- Both tools have the same philosophy:
 - Get information you are looking for in 2 - 3 clicks!
 - You can find utilization of any attached device in a simple graphical form understandable from technician to management level
 - Minimal effort for tools management
 - It can run on any Unix (Linux, AIX ...)
 - Both tools can run agent less
 - It is intended as a front-end tool

Introduction of LPAR2RRD

- Creates historical, future trends and nearly "real-time" CPU utilization graphs of LPAR's and shared CPU usage of IBM Power servers.
- Collects complete physical and logical configuration of all physical servers & LPAR's.
- It is agent less (getting all required data from management stations like the HMC/SDMC/FSM and IVM).
- It supports all kinds of logical partitions:
 - AIX / AS400 / Linux / VIOS

Introduction of LPAR2RRD

- It graphically presents:
 - CPU workload per partition
 - CPU per CPU Pool/physical server
 - Memory usage on physical server
- It natively supports following IBM technologies
 - CPU sharing
 - Live Partition Mobility
 - Active Memory Sharing
 - Active Memory Expansion
 - Capacity on Demand

Introduction of LPAR2RRD

- It graphically presents:
 - CPU workload per partition
 - CPU per CPU Pool/physical server
 - Memory usage on physical server
- It natively supports following IBM technologies
 - CPU sharing
 - Live Partition Mobility
 - Active Memory Sharing
 - Active Memory Expansion
 - Capacity on Demand

IT Capacity Planning

- What does it means?
- Does that save money?
- If so then how?
- Have you already have an IT capacity manager position and a tool at your company?

Capacity Planning

- Why do we need capacity planning?
 - To predict future bottlenecks
 - To predict future HW purchase with proper sizing
- How to start?
 - Need to have data
- Where to get data?
 - Need to have performance monitoring in place
- What to monitor?
 - CPU
 - Memory
 - Networking (LAN / SAN / WAN)
 - Storage

Capacity Planning - CPU

- Can be highly virtualized (IBM PowerVM, VMware ...)
 - Virtualized means shared between virtual partitions
- Lack of CPU power or bad virtualization design often cause bottlenecks
- CPUs are quite expensive
- A lot of SW products has licensing per CPU
 - Companies should care about number of CPU to not have to buy unnecessary SW licenses
 - Accurate capacity planning can directly save money spent for HW and SW
- Often happens that servers are fully CPU equipped and cannot be extended
 - Server upgrade and or LPAR migration can be long process
- **All above reasons indicate that CPU is most important from capacity planning point of view**

Capacity Planning - Memory

- Usually is not virtualized like CPU
 - There are exceptions like
 - IBM PowerVM Active Memory Sharing
 - VMware
- Consumption is not so dynamic as CPU
 - i.e. database memory consumption is fixed by cfg
- Mem size has usually no effect on SW licensing
- Easily extendable
 - Nowadays physical limits of memory on servers are usually high enough
 - Extend can be planned in days after a problem appears

Capacity Planning - Networking

- Bottleneck on the LAN / SAN is very rare
 - usually only backups are able regularly highly utilize LAN/SAN apart of some configuration mistakes or cluster interconnect
- When a problem appears then it is usually caused by bad infrastructure design
- Fix is usually fast and relatively cheap
 - new IO cards, network component, trunk ...
- From capacity planning view it is not major area
- Network usually don't cause a bottleneck!

Capacity Planning- Storage

- From capacity planning point of view
 - Storage capacity (size)
 - Storage throughput (IOPs, Bytes/sec)
- Storage capacity
 - Can be relatively simple retrieved and approximated to the future
 - When is no special requirement for throughput then new capacity is relatively cheap
 - Disk upgrade is usually quick operation
- Storage throughput
 - Storage throughput is often reason of bottlenecks
 - **Storage IO bottleneck** is usually though to prove or predict

CPU capacity planning


- Next part of presentation will be focused on CPU capacity planning on IBM Power Systems platform
- However the product is able to monitor VMware on Intel as well

CPU virtualization on IBM Power Systems

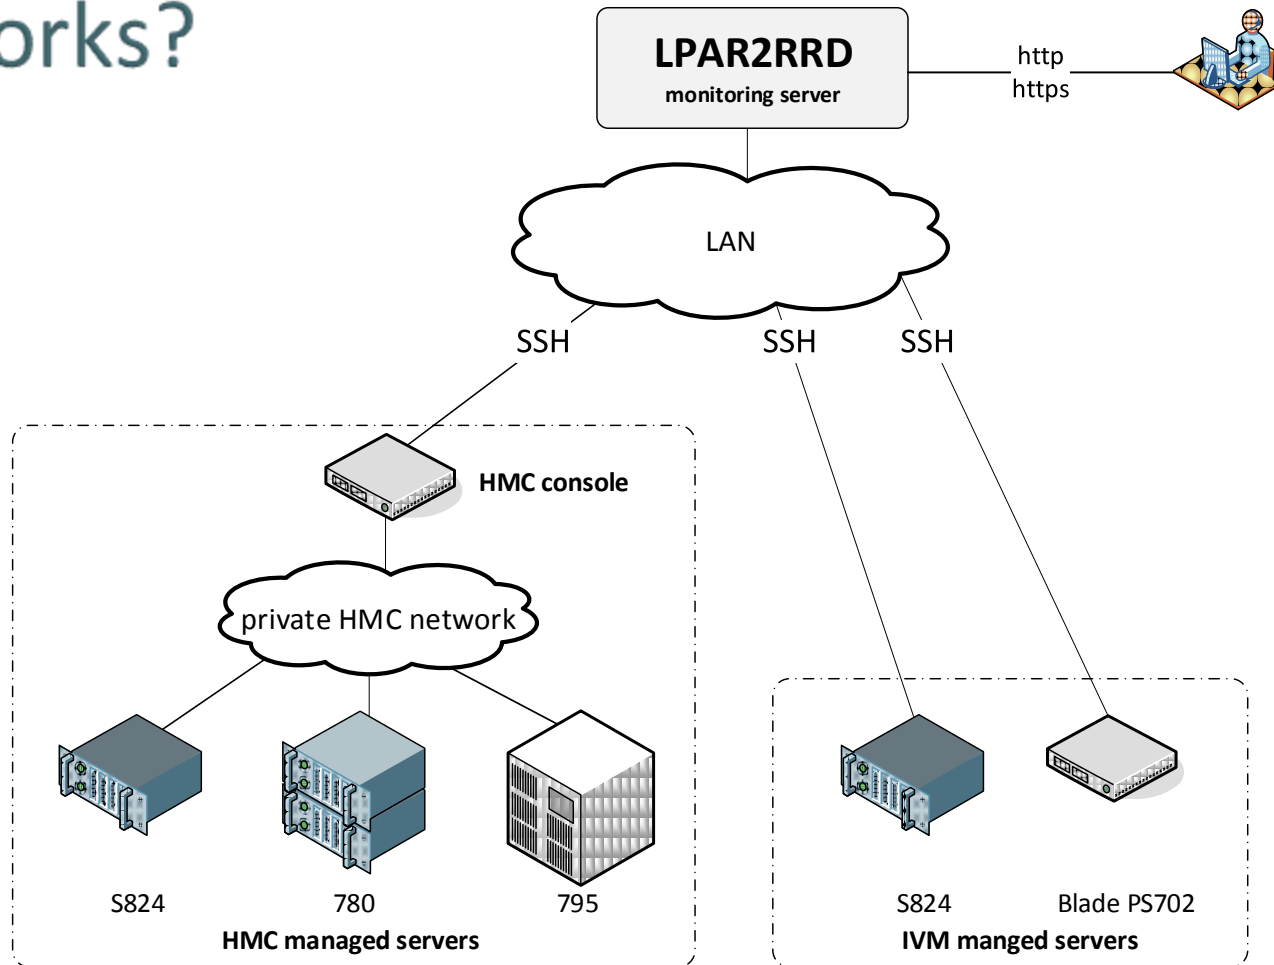
- The problem:
 - How to manage the CPU pools and sub pools on IBM Power Systems
 - Document actual setup
 - Knowledge of free resources
 - Determine CPU usage trends
 - Report consumption of resources per LPAR
 - Identify abnormal CPU usage
 - Manage Capacity
 - Compare CPU load on different server models for migration purposes

LPAR2RRD - highlights

- Light weight solution
- Real time and historical data
- Easy to implement and manage
- Cost effective (free)



How it works?




CPU trends

You can see, current trends in CPU consumption based on historical data – from last month to 1 year

More examples on: http://lpar2rrd.com/aix_ismseries_trends.htm

Or on demo.lpar2rrd.com



CPU Workload Estimator (CWE)

- It makes you CPU estimation based on your historical data
- Useful for migration / consolidation
- Estimation based on rPerf (AIX) or CPW (iOS) IBM benchmark
- Appropriate SMT level could be selected

CWE

- Just select current physical servers or LPARs considered for migration
- Chose if you will consolidate on existing server or plan for new one
- And generate report, to see, if you fit in desired box
- You can very easily create many different scenarios to chose the right solution for you

CWE – example 1

- Consolidation on existing server
- In this example we try to fit LPAR from one existing server to another existing server
- Run CWE, in left section select LPAR, target server is existing and select destination pool
- Then press Generate Report button

CPU Workload Estimator

From to

Graph resolution x Y-axis

LPAR(s) for migration

Target server:

Server | LPAR all

Filter...

- ASRV11
 - ASRV11LPAR10 [Pcas11]
 - ASRV11LPAR11 [Tcas11]
 - ASRV11LPAR12 [Acas11]
 - ASRV11LPAR13 [Pras11]
 - ASRV11LPAR14 [Tras11]
 - ASRV11LPAR15 [Aras11]
 - ASRV11LPAR16 [Peas11]
 - ASRV11LPAR17 [Teas11]
 - ASRV11LPAR18 [Aeas11]
 - ASRV11LPAR19 [P1ias11]
 - ASRV11LPAR20 [T1ias11]
 - ASRV11LPAR21 [A1ias11]
 - ASRV11LPAR22 [P2ias11]
 - ASRV11LPAR23 [T2ias11]
 - ASRV11LPAR24 [A2ias11]
 - ASRV11LPAR4
 - ASRV11LPAR5 [sse11]
 - ASRV11LPAR6 [ase11]
 - ASRV11LPAR7
 - ASRV11LPAR8

Server | Pool

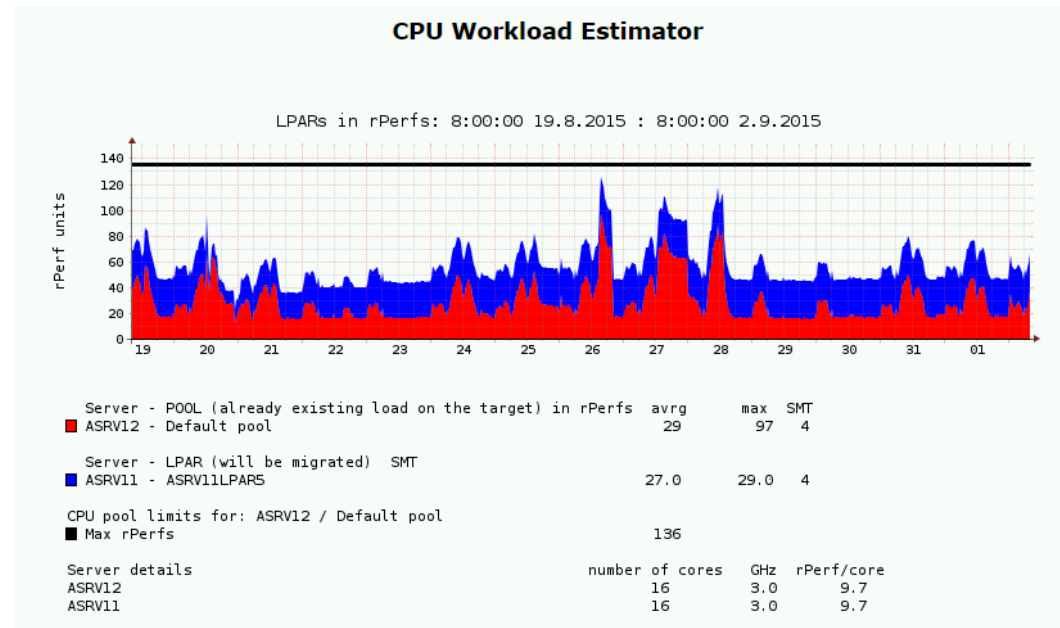
- ASRV11
- ASRV12
 - CPU pool
- BSRV21
- BSRV22

New server type

| Platform/Model | Type | CPU | GHz |
|----------------|------|-------|------|
| Power8 | | | |
| 8408-E8E1 | E850 | P8/16 | 3.72 |
| 8408-E8E1 | E850 | P8/24 | 3.72 |
| 8408-E8E1 | E850 | P8/32 | 3.72 |
| 8408-E8E2 | E850 | P8/20 | 3.35 |
| 8408-E8E2 | E850 | P8/30 | 3.35 |
| 8408-E8E2 | E850 | P8/40 | 3.35 |
| 8408-E8E3 | E850 | P8/24 | 3.02 |
| 8408-E8E3 | E850 | P8/36 | 3.02 |
| 8408-E8E3 | E850 | P8/48 | 3.02 |
| 9119-MME | E870 | P8/8 | 4.02 |
| 9119-MME | E870 | P8/12 | 4.02 |
| 9119-MME | E870 | P8/16 | 4.02 |
| 9119-MME | E870 | P8/20 | 4.02 |
| 9119-MME | E870 | P8/24 | 4.02 |
| 9119-MME | E870 | P8/28 | 4.02 |
| 9119-MME | E870 | P8/32 | 4.02 |
| 9119-MME | E870 | P8/64 | 4.02 |
| 9119-MME | E870 | P8/40 | 4.19 |

CWE – example 1

- As you can see, this LPAR (blue part of graph) will fit to current workload (red part of graph) on this particular server



CWE – example 2

- Server replacement
- In this example we try to fit workload from one existing physical server to brand new physical server
- Run CWE, in left section select complete physical server, target server is new and select required model
- Then press Generate Report button

CPU Workload Estimator

From to
 Graph resolution x Y-axis

LPAR(s) for migration Target server:

Server | LPAR all

Filter...

- ASRV11
- ASRV12
- BSRV21
- BSRV22

➔

Server | Pool


- ASRV11
- ASRV12
- CPU pool
- BSRV21
- BSRV22

New server type

| | | | |
|----------|------|--------|------|
| 9119-MHE | E880 | P8/40 | 4.00 |
| 9119-MHE | E880 | P8/44 | 4.00 |
| 9119-MHE | E880 | P8/48 | 4.00 |
| 9119-MHE | E880 | P8/96 | 4.00 |
| 9119-MHE | E880 | P8/144 | 4.00 |
| 9119-MHE | E880 | P8/192 | 4.00 |
| 8286-41A | S814 | P8/4 | 3.02 |
| 8286-41A | S814 | P8/6 | 3.02 |
| 8286-41A | S814 | P8/8 | 3.72 |
| 8284-22A | S822 | P8/6 | 3.89 |
| 8284-22A | S822 | P8/8 | 4.15 |
| 8284-22A | S822 | P8/10 | 3.42 |
| 8284-22A | S822 | P8/12 | 3.89 |
| 8284-22A | S822 | P8/16 | 4.15 |
| 8284-22A | S822 | P8/20 | 3.42 |
| 8286-42A | S824 | P8/6 | 3.89 |
| 8286-42A | S824 | P8/8 | 4.15 |
| 8286-42A | S824 | P8/12 | 3.89 |
| 8286-42A | S824 | P8/16 | 4.15 |
| 8286-42A | S824 | P8/24 | 3.52 |

CWE – example 2

- As you can see, this physical server will fit to this new server
- Black lines in graph represents different SMT values
- All LPARs in this example runs at SMT4 which is represented by 2nd line from top
- In this case we have approximately 20rPerf of spare capacity



CWE – example 3

- Server replacement and consolidation
- In this example we try to fit workload from all existing physical servers to one brand new physical server
- Run CWE, in left section select all complete physical servers, target server is new and select required model
- Then press Generate Report button

CPU Workload Estimator

From to
Graph resolution x Y-axis

LPAR(s) for migration [HMC Info](#) Target server:

Server | LPAR all

Filter...

- ASRV11
- ASRV12
- BSRV21
- BSRV22

Server | Pool

- ASRV11
- ASRV12
- CPU pool
- BSRV21
- BSRV22

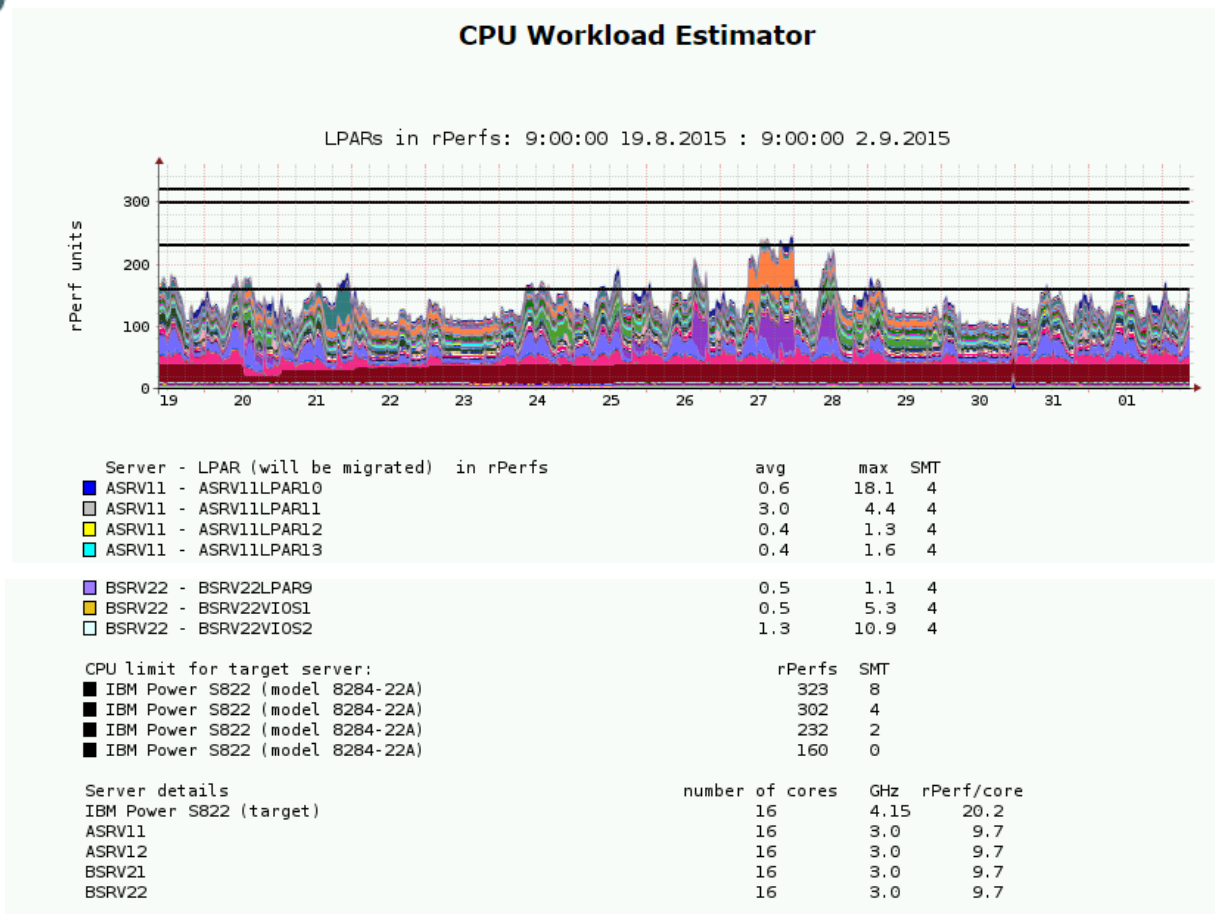
New server type

| | | | |
|----------|------|-------|------|
| 8286-41A | S814 | P8/4 | 3.02 |
| 8286-41A | S814 | P8/6 | 3.02 |
| 8286-41A | S814 | P8/8 | 3.72 |
| 8284-22A | S822 | P8/6 | 3.89 |
| 8284-22A | S822 | P8/8 | 4.15 |
| 8284-22A | S822 | P8/10 | 3.42 |
| 8284-22A | S822 | P8/12 | 3.89 |
| 8284-22A | S822 | P8/16 | 4.15 |
| 8284-22A | S822 | P8/20 | 3.42 |
| 8286-42A | S824 | P8/6 | 3.89 |
| 8286-42A | S824 | P8/8 | 4.15 |
| 8286-42A | S824 | P8/12 | 3.89 |
| 8286-42A | S824 | P8/16 | 4.15 |
| 8286-42A | S824 | P8/24 | 3.52 |

- Power7+
- Power7
- Power6+
- Power6
- Power5+
- Power5

CWE – example 3

- As you can see, this workload will fit to this new server
- Black lines in graph represents different SMT values
- All LPARs in this example runs at SMT4 which is represented by 2nd line from top
- In this case we have approximately 65rPerf of spare capacity



CWE - Customer Benefits

- Automated capacity planning and performance analysis
- Provides the information needed to help:
 - productively and cost effectively manage their growth and performance
- Ensure uninterrupted availability of system
- The customer be prepared for IT challenges / opportunities
- Maximize return on current and future hardware investments with more effective utilization
- Easily understand current workload and plan for future requirements

Business model of LPAR2RRD tool

- Product is free
- Support is for fee
- Support levels
 - Standard
 - Premium
- Some features are available only for customers under support

Business model of LPAR2RRD tool

- Product is free
- Support is for fee
- Support levels
 - Standard
 - Premium
- Some features are available only for customers under support

Support for LPAR2RRD

- You might optionally order support and you get:
 - Standard at least:
 - priority in bug fixing, installation support, etc
 - Premium at least:
 - defined SLA's (1 working day for critical issues)
 - new functionality implementation as per request
 - report customizations
 - prolong data retention time (received data is stored for 60days then its being averaged)
 - with support you get additional functionality
 - CPU Workload Estimator based on rPerf and CPW benchmarks
 - Live Partition Mobility support
 - Custom Groups (unlimited number of LPARs in a group)

LPAR2RRD - Installation Requirements

- **All can be setup within 1 hour!**
- LPAR2RRD can run on any Unix OS
- Disk space requirements are about 4MB per monitored LPAR
- Required SW
 - Any web server
 - SSH
 - RRDTools
 - Perl

Introduction of STOR2RRD

- Free performance monitoring tool for storages with following features:
 - Create storage utilization graphs for monitored storages
 - Create historical and nearly on-line utilization graphs
- It graphically presents:
 - IO rate, data throughput, response times
 - front-end and back-end data
 - Ports, Pools, Arrays, Ranks, Mdisks, Volumes, Drives
 - Host aggregated graphs

Introduction of STOR2RRD

- It comes with the same philosophy as its sister product LPAR2RRD:
 - **Get information you are looking for in 2 - 3 clicks!**
- You can find utilization of any attached storage in a simple graphical form understandable from technician to management level
- Minimized tool management
- It can run on any Unix (Linux, AIX ...)

Introduction of STOR2RRD

- Following storages are supported (as of Aug 2015)
 - IBM System Storage DS8000 series
 - IBM System Storage DS6800 series
 - IBM Storwize family
 - IBM SAN Volume Controller (SVC)
 - IBM XIV
 - IBM DS3000, DS4000 and DS5000
 - IBM DCS3700
 - NetApp E-Series
 - Brocade SAN switches

Introduction of STOR2RRD

- Metrics
 - IO rate [IO per sec]
 - data throughput [MB/sec]
 - response times [ms]
 - cache hits and usage
 - PPRC (DS8000), IP replication (Storwize/SVC)
 - Pool capacity usage [GB]
 - CPU utilization, compression usage (Storwize/SVC)
- **All read/write (in/out) separately**
- **Storage front-end and back-end stats**

STOR2RRD

- Monitored logical and physical components
 - Port (FC, SAS, PCIe, FCoE, iSCSI, IP replication)
 - Pool and Array
 - RANK and Mdisk
 - Volume (LUN)
 - Drive (Storwize and SVC)
 - Hosts
- feature matrix:
www.stor2rrd.com/support_matrix.htm

STOR2RRD - Forecasted enhancements

- Support of further storages
 - EMC, Hitachi and NetApp midrange storages have priority
- Features for capacity tracking and planning
- Integration with LPAR2RRD
- Automated documentation of monitored storages
- Alerting for:
 - Alerting feature going to be released in the mid of Q3 2015

STOR2RRD - example

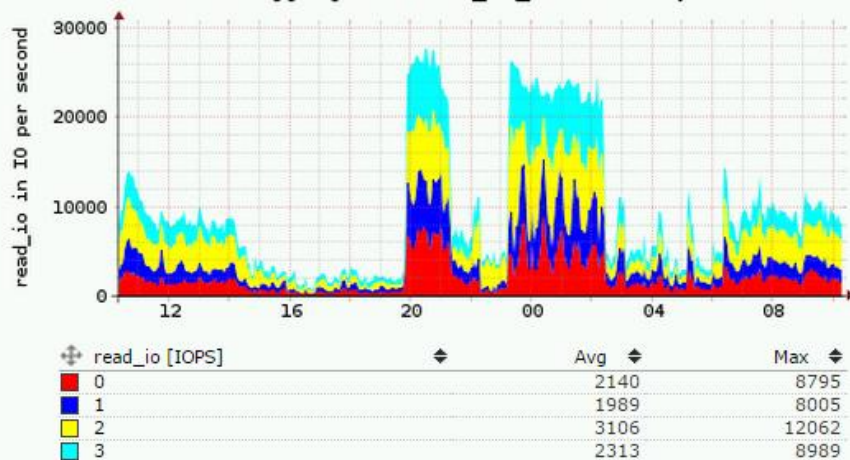


STOR2RRD - example

POOL aggregated sum_io: last day



POOL aggregated read_io_b: last day



References

- **LPAR2RRD home page**
 - <http://www.lpar2rrd.com>
 - **Demo site**
 - <http://demo.lpar2rrd.com/>
- **STOR2RRD home page**
 - <http://www.stor2rrd.com>
 - **Demo site**
 - <http://demo.stor2rrd.com/>