IBM Systems Lab Services

Infrastructure expertise to help you build the foundation of a smart enterprise

Lab Services Tools

Turgut Genc - Senior Consultant

IBM Systems Lab Services, EMEA Geo leader:

Power Enterprise Pools – Lab Services tools – Migration and Performance workshops.

E-mail : TurgutGenc@uk.ibm.com
Twitter : @turgutge
Slack : @TurgutGenc

ibm.biz/IBMLabServices
ibmsls@us.ibm.com

Agenda

Who are we?
Lab Services Tools overview
Lab Services Tools in details
Demonstrations

- Demo 1: LPM & SRR Automation Toolkit
- Demo 2: PowerDraw
- Demo 3: Capacity on Demand Dashboard

Questions

IBM Systems Lab Services Organization

IBM Systems Lab Services has the infrastructure expertise to help you build the foundation of a smart enterprise.

From servers and mainframes to storage systems and software, Lab Services helps you deploy the building blocks of a next-generation IT infrastructure that empowers your business.

We accelerate adoption of systems and storage with infrastructure services for 8,000+ clients worldwide annually.

Lab Services also promotes best practices and shares experience with clients and ecosystem partners via IBM TechU events.

IBM Systems Lab Services Consultants

Our Lab Services consultants perform infrastructure services for clients on site, offering deep technical expertise, valuable tools and successful methodologies.

Our services are designed to help clients solve business challenges, gain new skills and apply best practices.

Lab Services offers a wide range of infrastructure services for:

- IBM Power Systems and Cognitive
- IBM Z and IBM LinuxONE
- IBM Storage and Software Defined Infrastructure

Lab Services has a global presence and can deploy experienced consultants around the world.

Lab Services Tools

Lab Services Tools Overview

Who

- IBM Systems Lab Services consultants develop these tools.

Why

- We address clients' needs while simplifying and automating the process with a tool
- While providing enhancements to existing capabilities
- Useful for Migration, Maintenance, Monitoring and Documentation purposes

How

- We only deliver these tools with a client engagement.
- Most of the tools are "free with an engagement"

  => Exception:LPM & SRR Automation Toolkit require purchase of an Enterprise License.
- Delivery engagement can be covered via Power to Cloud Rewards program, direct contract, etc.

  => Work with our Opportunity Managers in your location

Lab Services Tools: Migration

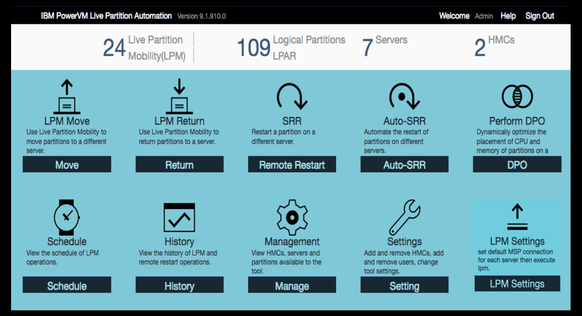

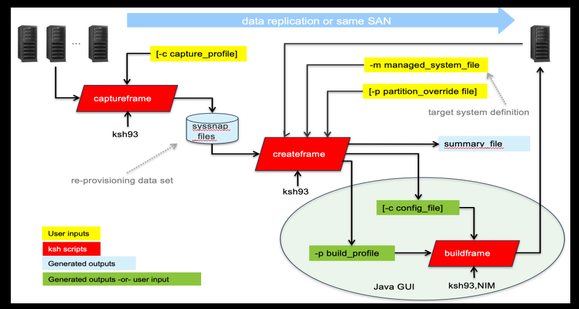

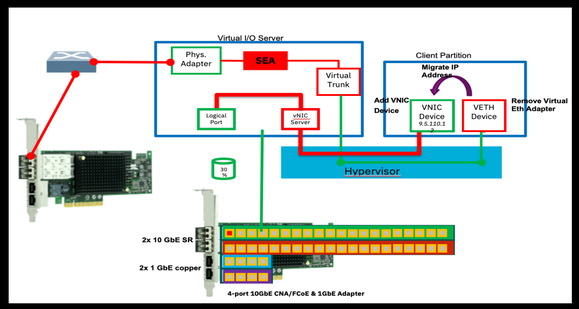

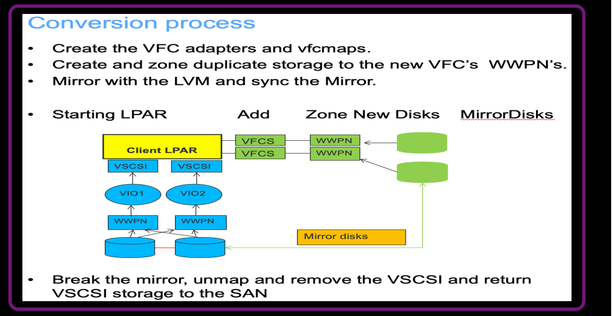

LPM & SRR Automation Tool v9

Design, automate and accelerate mobility operations for maintenance, migration and outages

PowerVM Provisioning Toolkit v8

Provision bare metal systems, capture/build from existing systems, orchestrate LPAR Migrations and decommissioning

SR-IOV Migration Tool

An assistance tool for live-migration from virtual Ethernet to SR-IOV devices

VSCSI-to-NPIV Migration Tool

Convert a VSCSI partition to NPIV with no downtime

Lab Services Tools: Monitoring

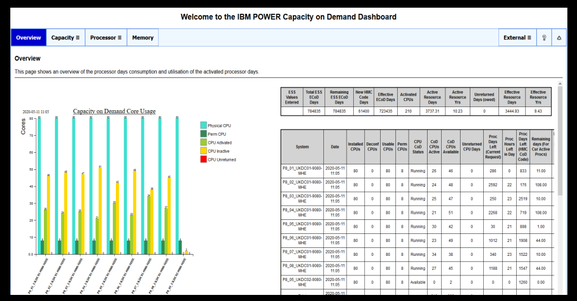

Capacity on Demand Dashboard

Monitoring CoD utilisation data to provide insight on Capacity usage, burn rate for each system, and a comparison of resource assignments

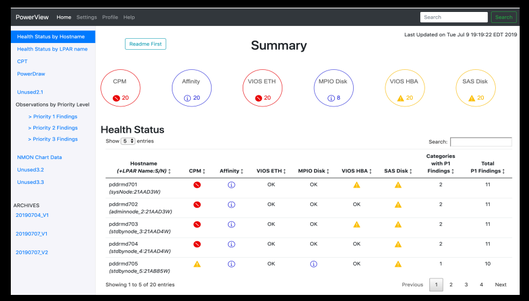

PowerView

An automated health check assessment toolset, AIX support staff could deploy for ongoing systems compliance validation

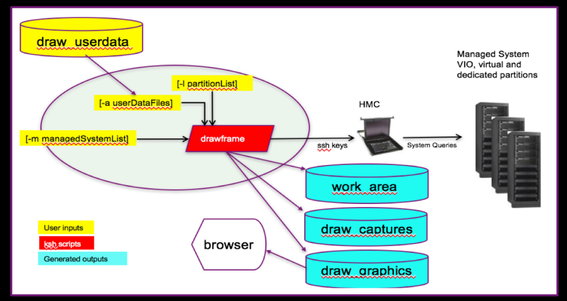

PowerDraw

Gather and display key configuration information of systems. Ability to perform problem determination using a graphical approach

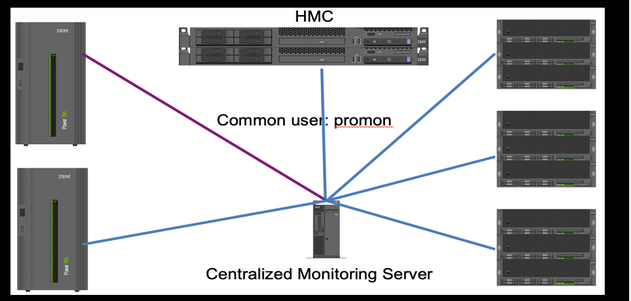

Proactive Monitoring w/ Promon

All systems tend toward chaos, keep environment tidy using this strategic monitoring tool

Lab Services Tools: Capacity Planning

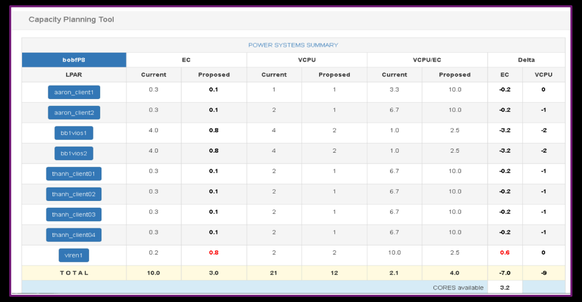

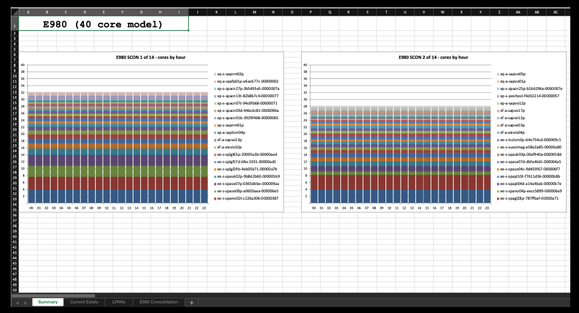

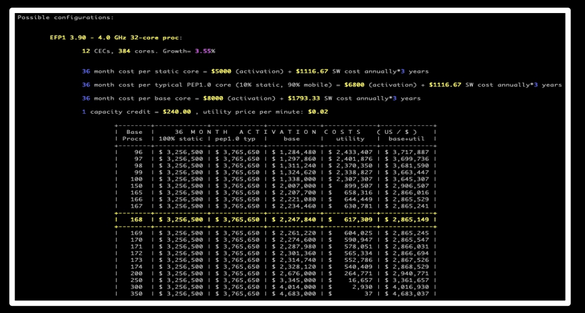

Capacity Planning Tool (CPT)

Use customer's utilization data and suggests right-sizing partitions' CPU

CPT-SCONN

Server Consolidation – Correct sizing for migration from one Power platform to another using rPerf scaling

CPT-MCM

Metered Capacity Modeling - Pools 2.0 Value Assessment Tool. Find the optimal number of base licenses for Pools 2.0

LPM and SRR Automation Toolkit v9

LPM and SRR Automation Tool

Design, Automate, Rollback

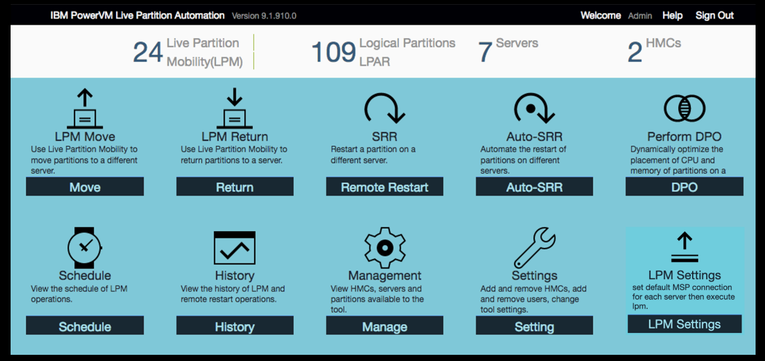

Live Partition Mobility (LPM)

Design for maintenance and migration

- Build a LPM plan for a maintenance window with control over VIO, HBA mappings, etc.

Automate and accelerate mobility action

- Schedule automated LPM operations or quickly move one or many partitions to another server in as few as 4 clicks with an easy-to-use GUI

Rollback simply to original server

- Return the partition/s back in a few as 4 clicks to the original server with the original HBA and Virtual slot ID mappings

Simplified Remote Restart (SRR)

Design for unplanned outages

- Build a SRR plan ready to execute in the event of an unplanned outage

Automate and accelerate mobility action

- Use a GUI to quickly SRR many or all the partitions to one or more destination servers

Rollback simply to original server

- Once the outage has been resolved / repaired, move all the partitions back to the original server with just a few clicks

LPM and SRR Automation Tool

Design, Automate, Rollback

Supports Power7 and Power8 and Power9 servers

Can be installed on AIX and Windows and Linux platforms. It is packaged as a zip file and contains all the code/packages needed to run.

Only communicates to the HMCs in your environment via ssh issuing HMC CLI commands. There is no need for agents or access to the VIOS or client partitions.

Only takes minutes to install the tool and connect to the HMCs and start using the tool.

LPM and SRR Automation Tool

Design, Automate, Rollback

Move many partitions at a time and the tool remembers where they need to go back to

Customize partition moves to your environment without becoming an HMC Command Line expert

Plan your partitions moves days in advance and Click and Go during maintenance window

Move partitions back to their original servers with previous as is configuration when needed.

LPM and SRR Automation Tool

Design, Automate, Rollback

Support for new LPM and SRR features

Provide GUI for SRR operations

Automate SRR operations

Ability to roll back after a SRR operation

LPM and SRR Automation Tool

Design, Automate, Rollback

Support for LDAP Authentication

Ability to define default MSPs (VIOS and IP Address) per server

Ability to integrate operations with PEP 1.0

Ability for pre-LPM, post-LPM, pre-SRR and pos-SRR scripting

Using the tool – login to tool

Launch a browser and point to the server where the tool is installed

i.e. https://<server where you installed the tool>:8443/lpm

(make sure you use this complete syntax as some browsers don't like shortened URLs)

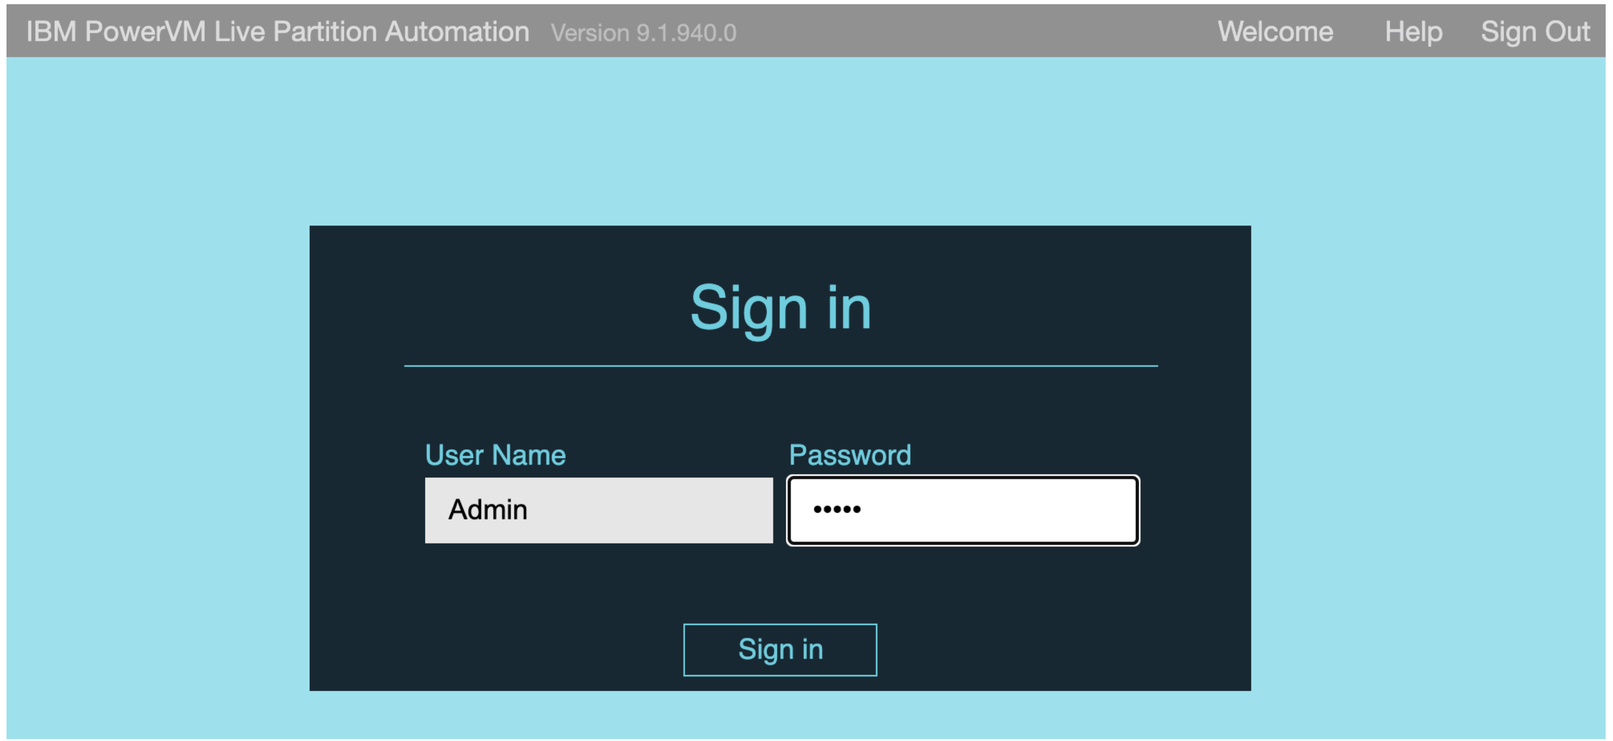

Home Screen – V9.9.940 release

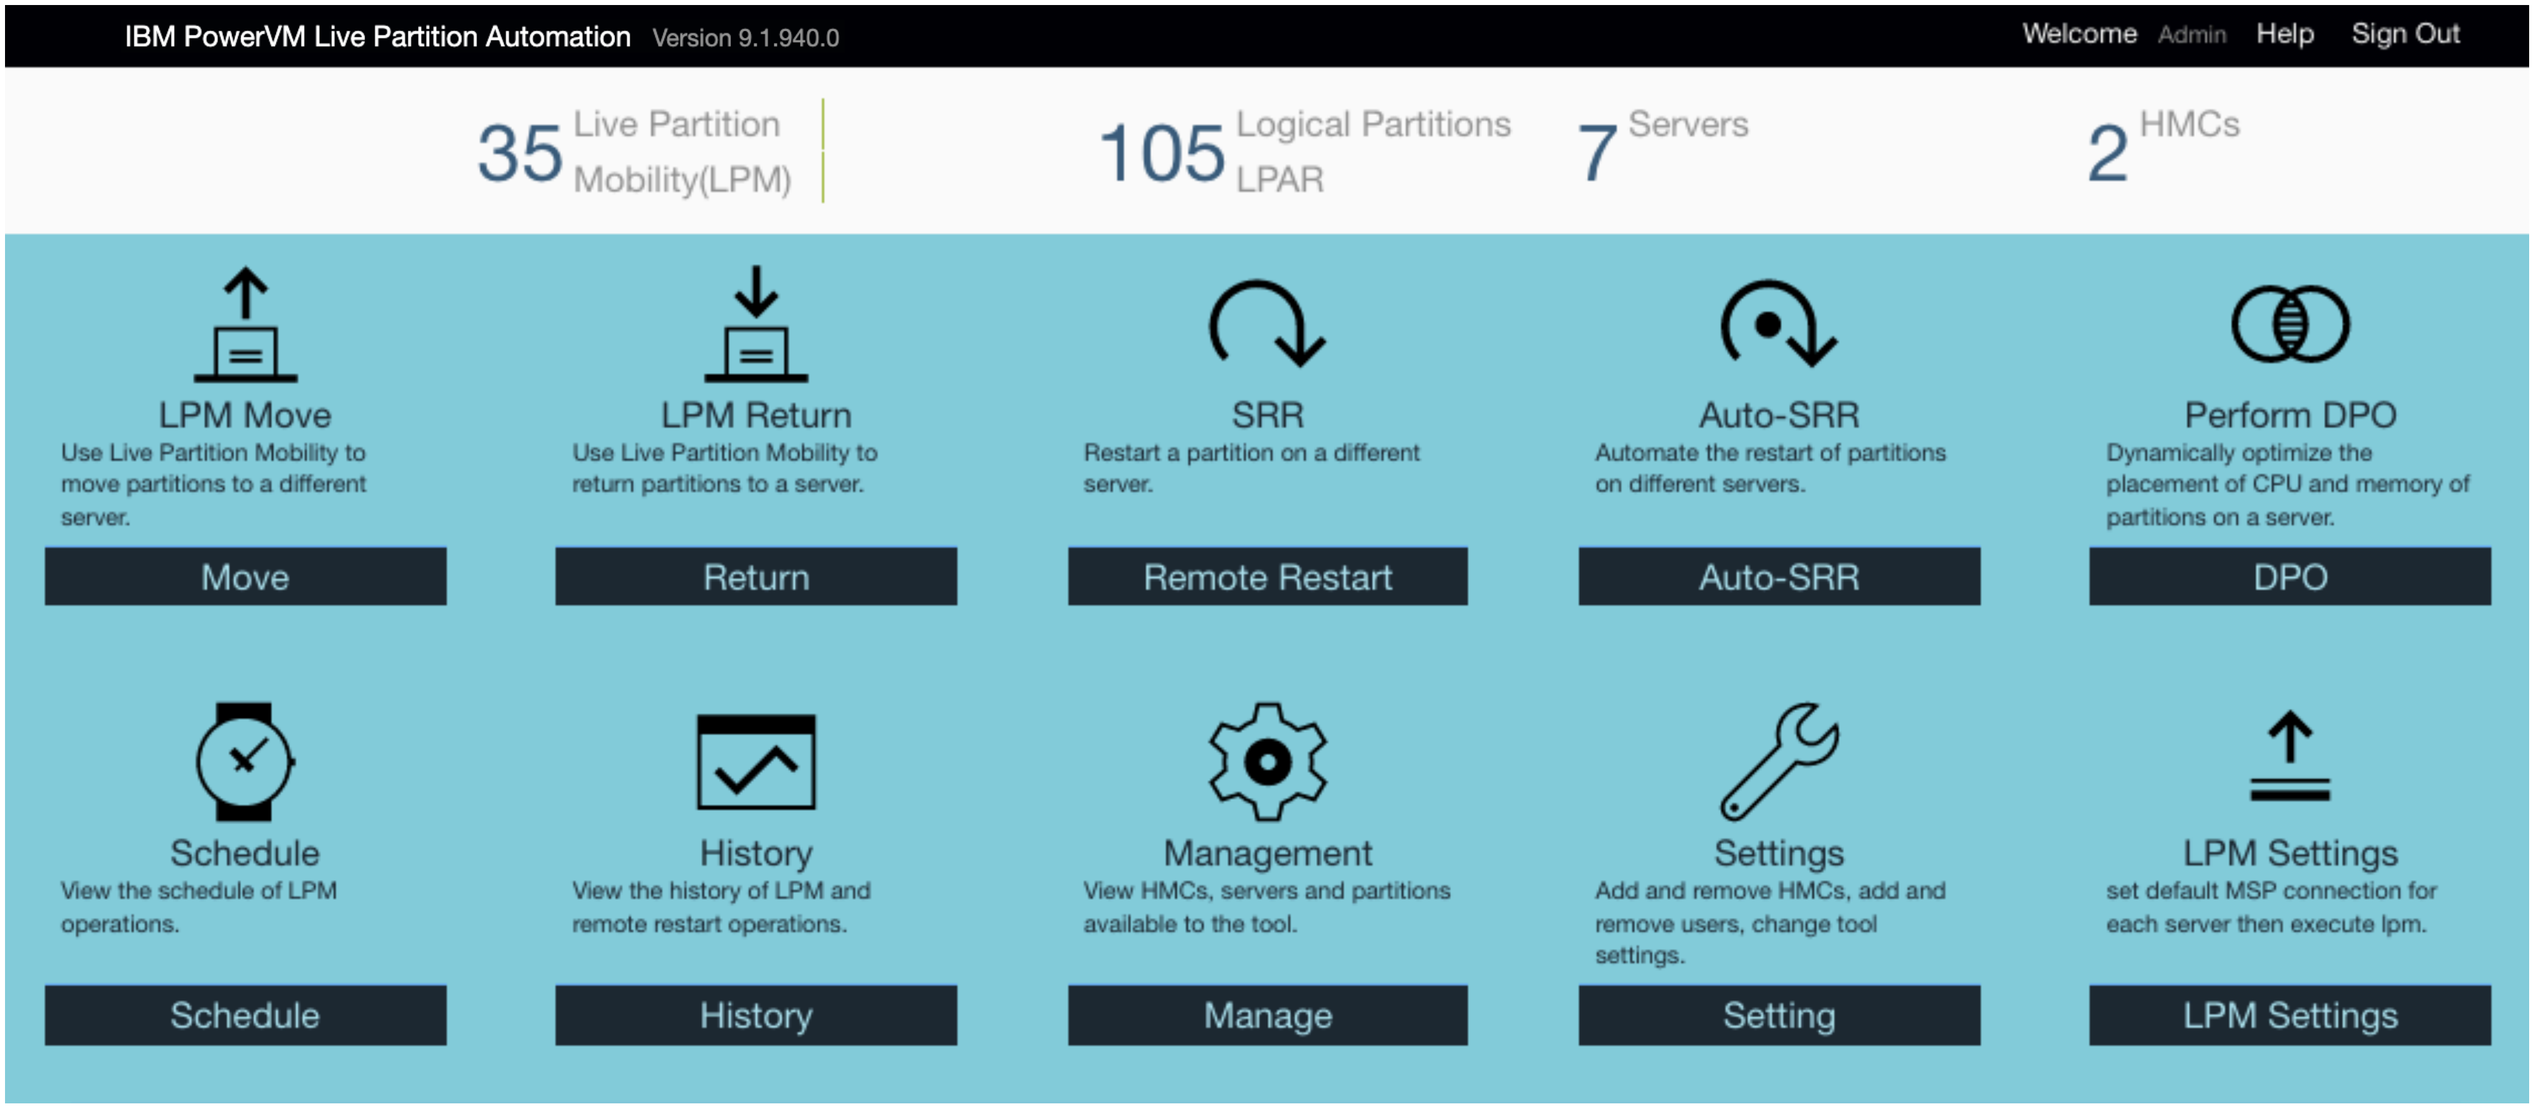

LPM Move: One Source & Multiple Destinations

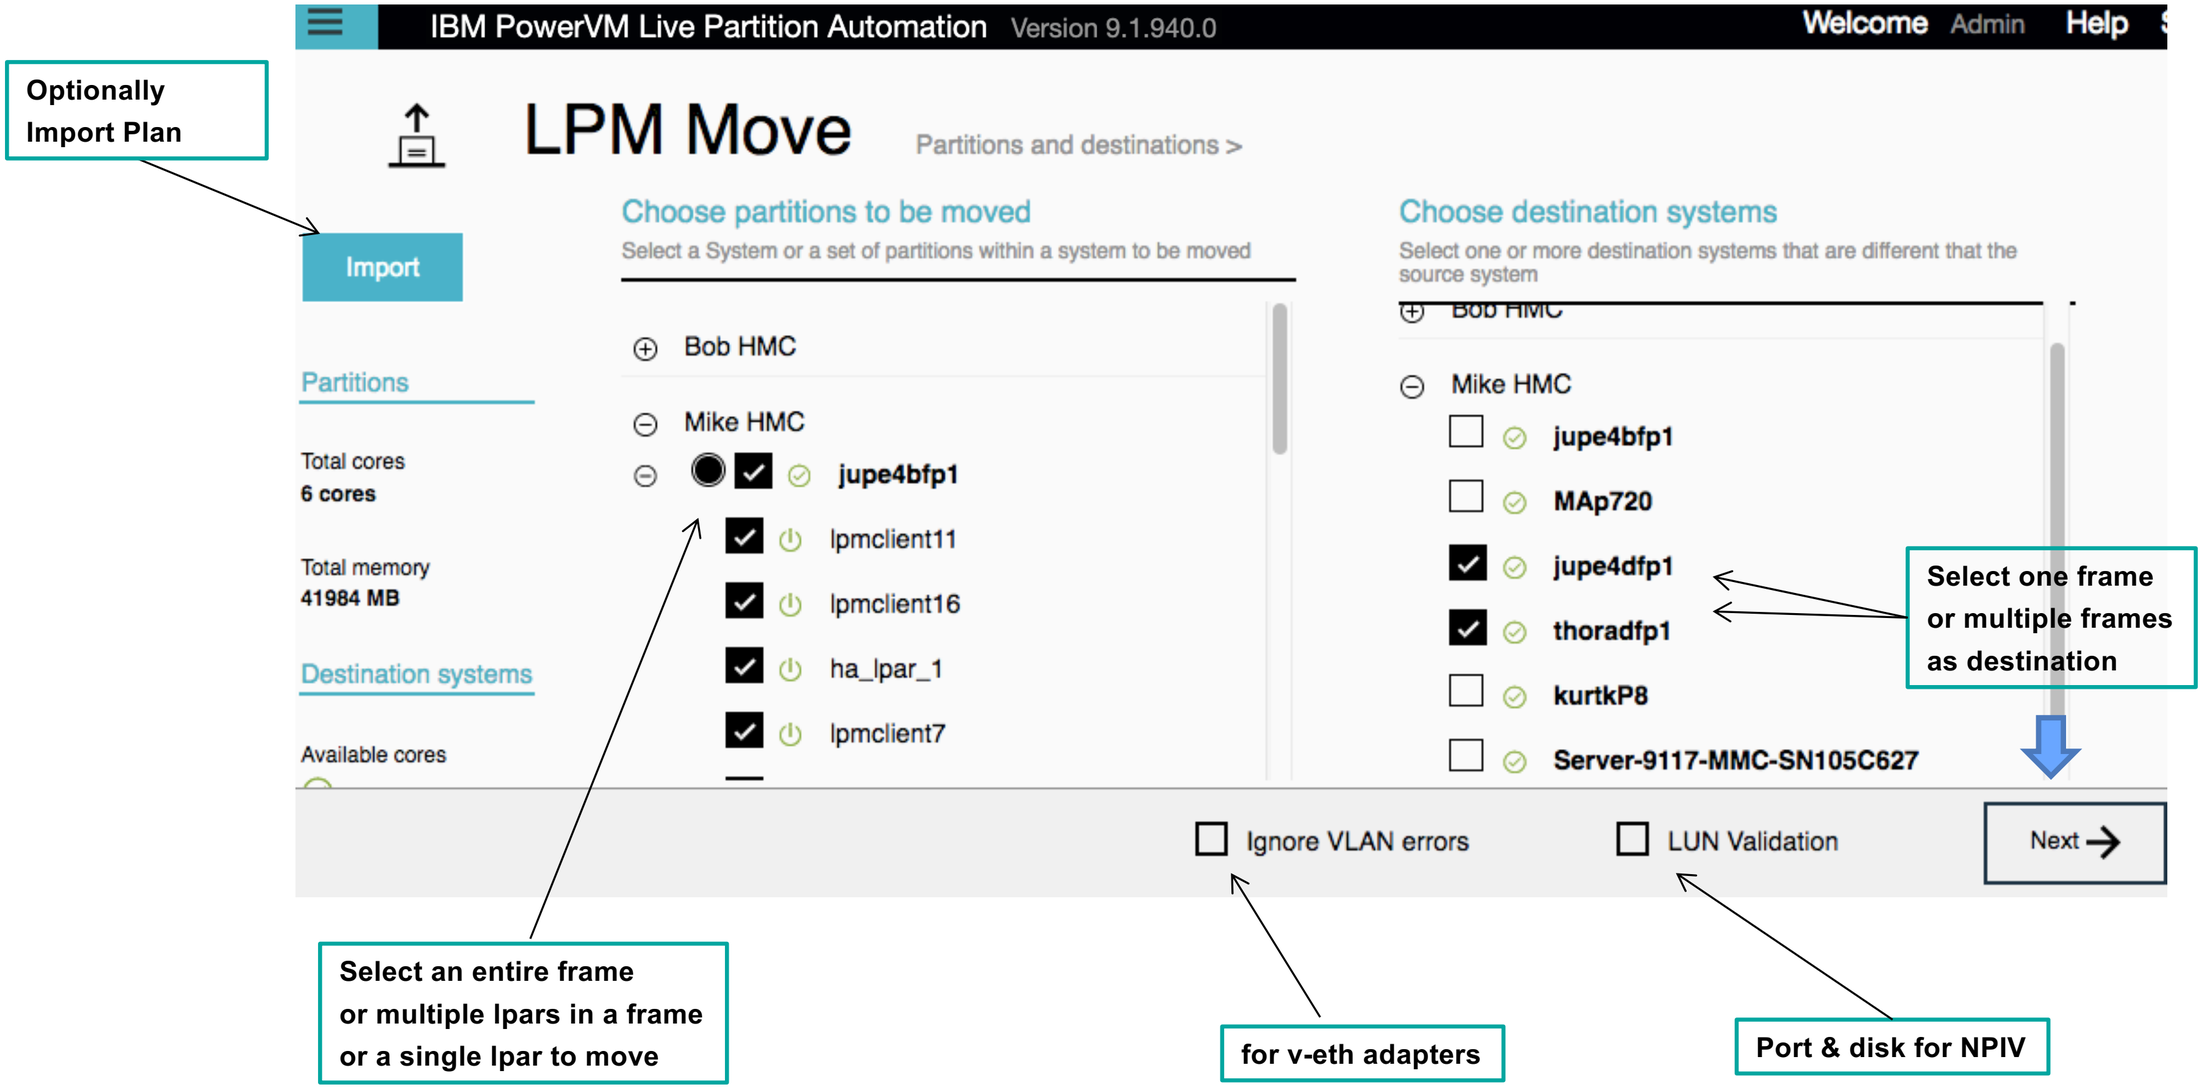

LPM Validation In Progress

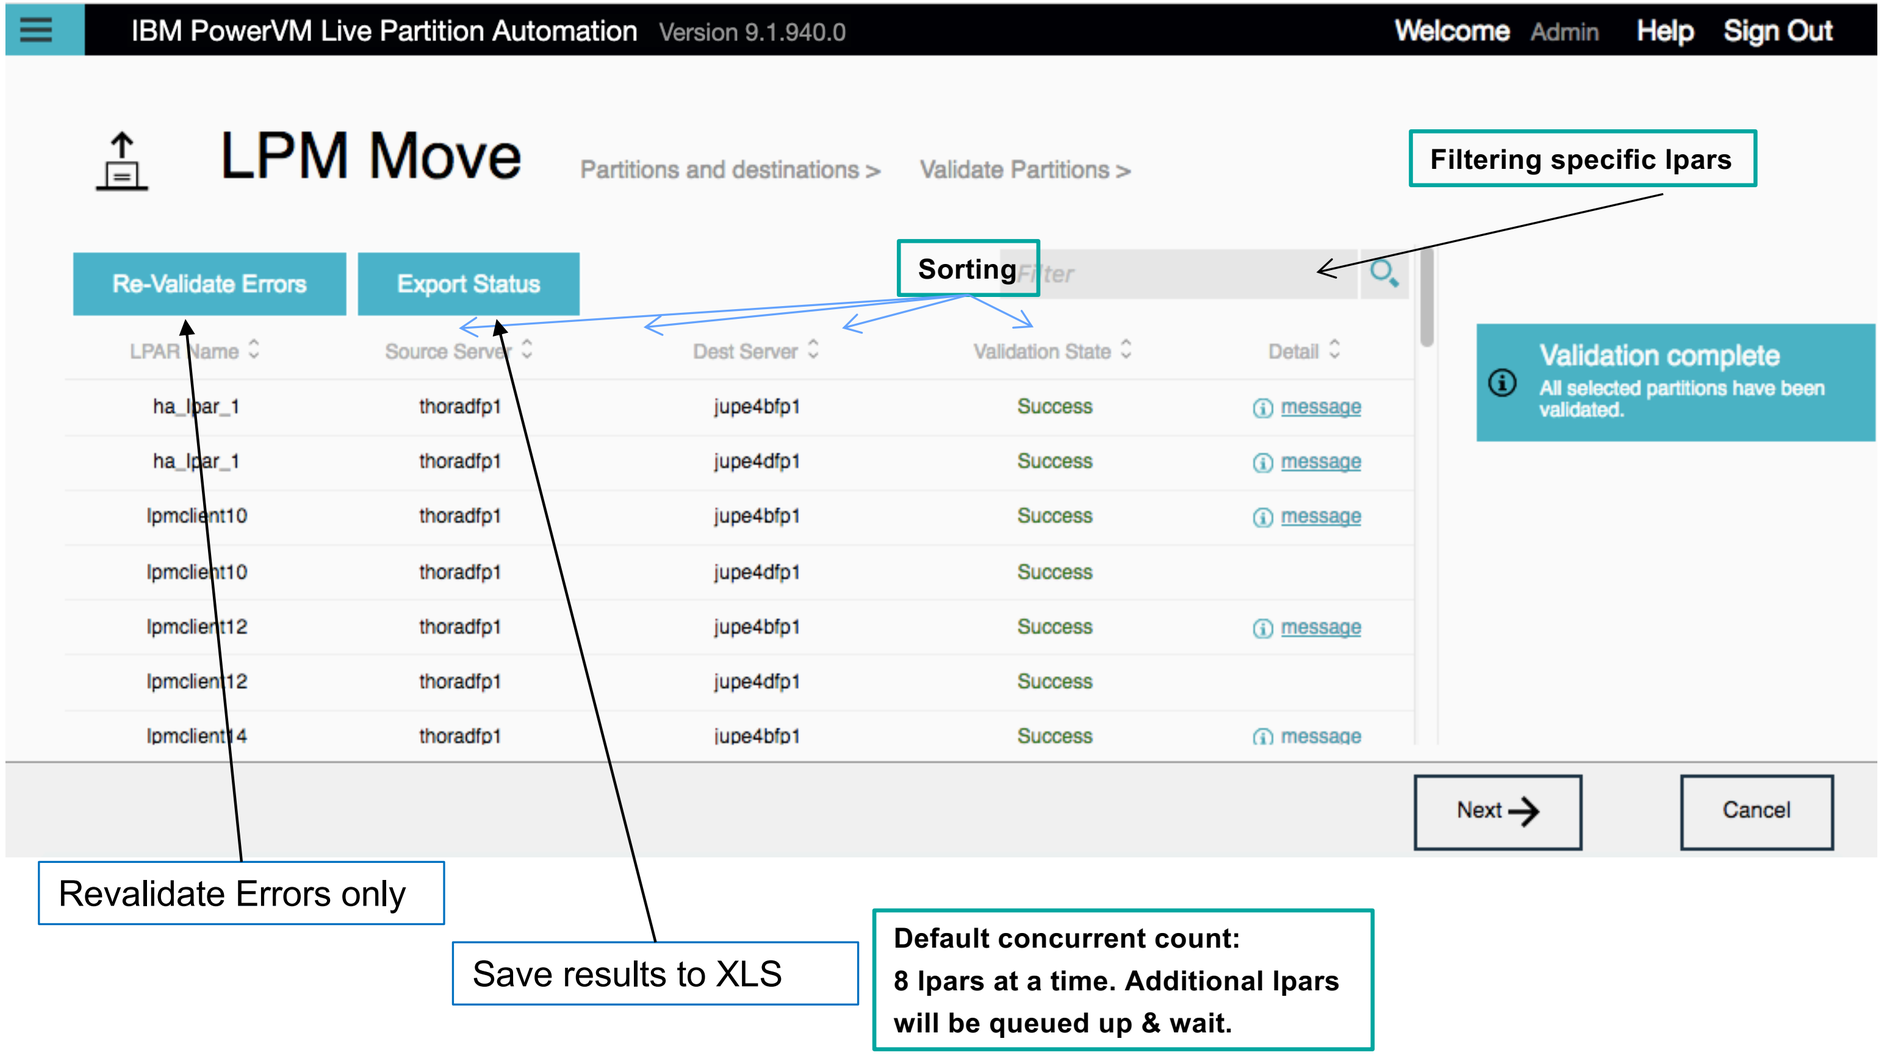

LPM Move – Placement

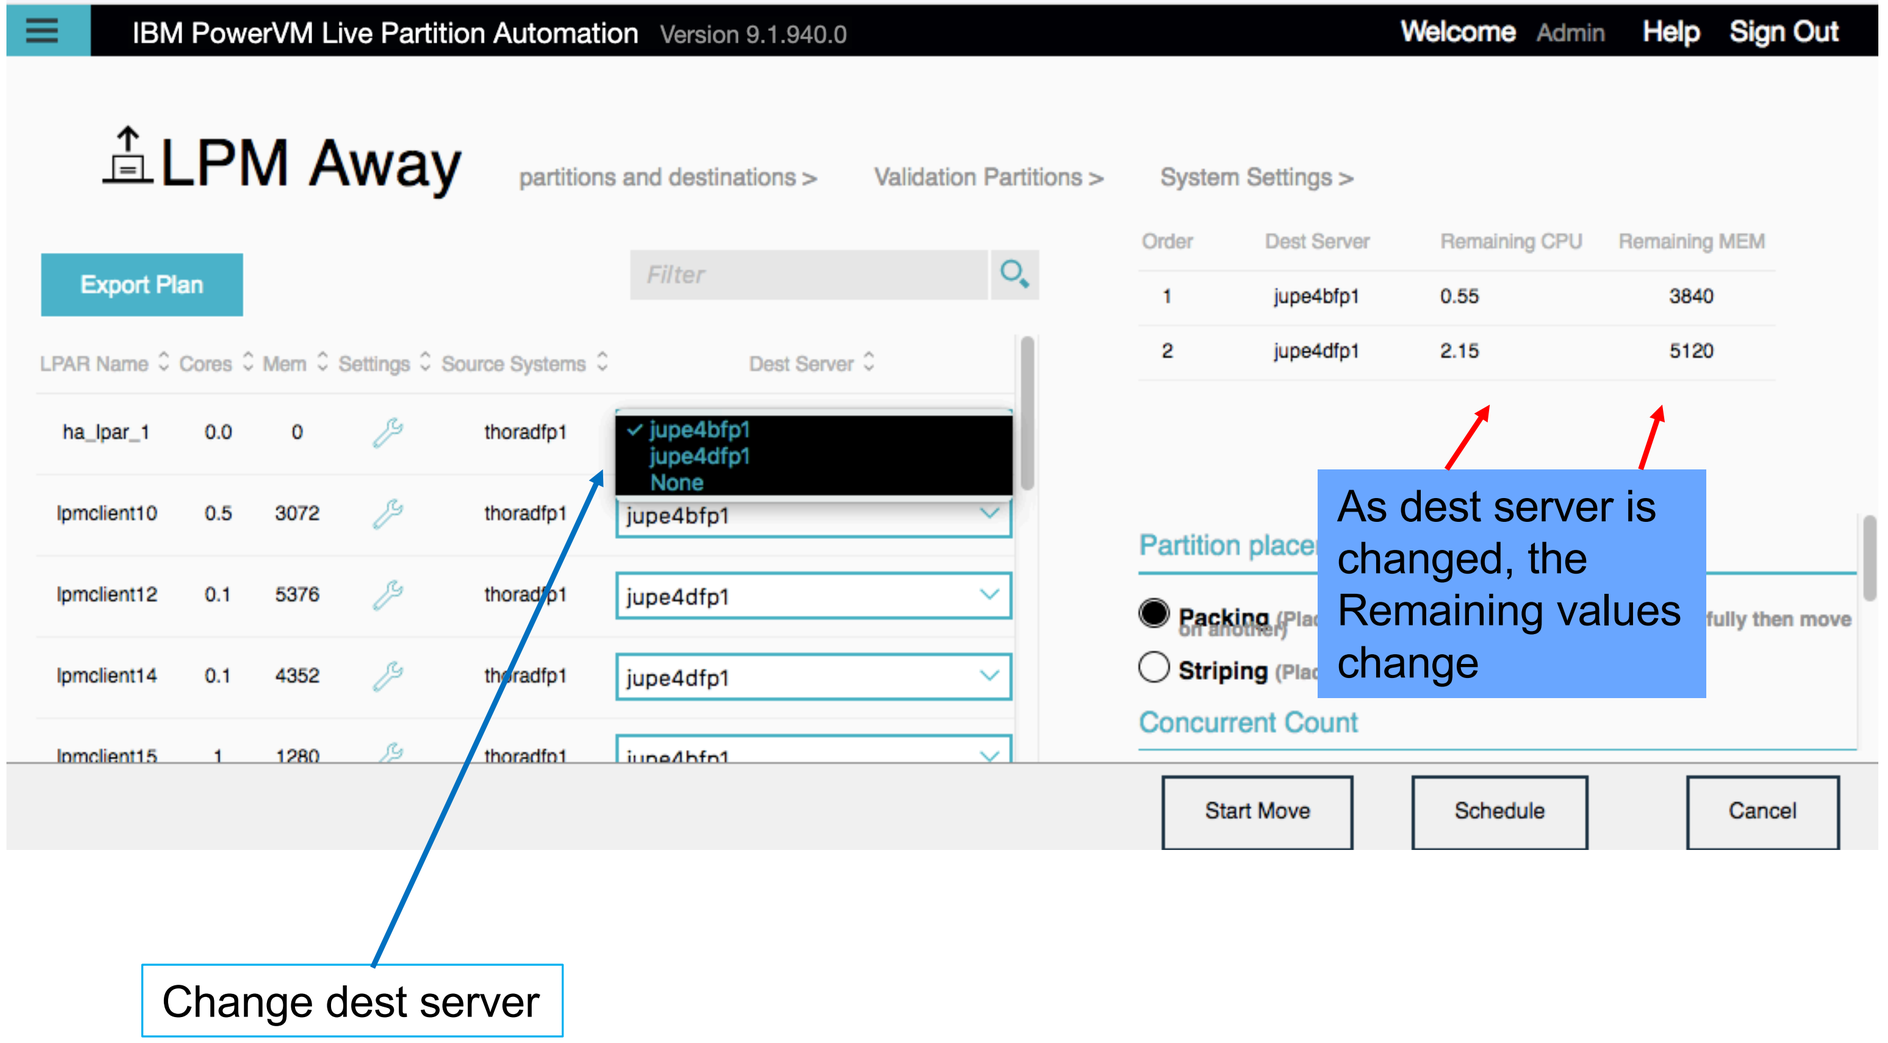

LPM Move – migration in progress

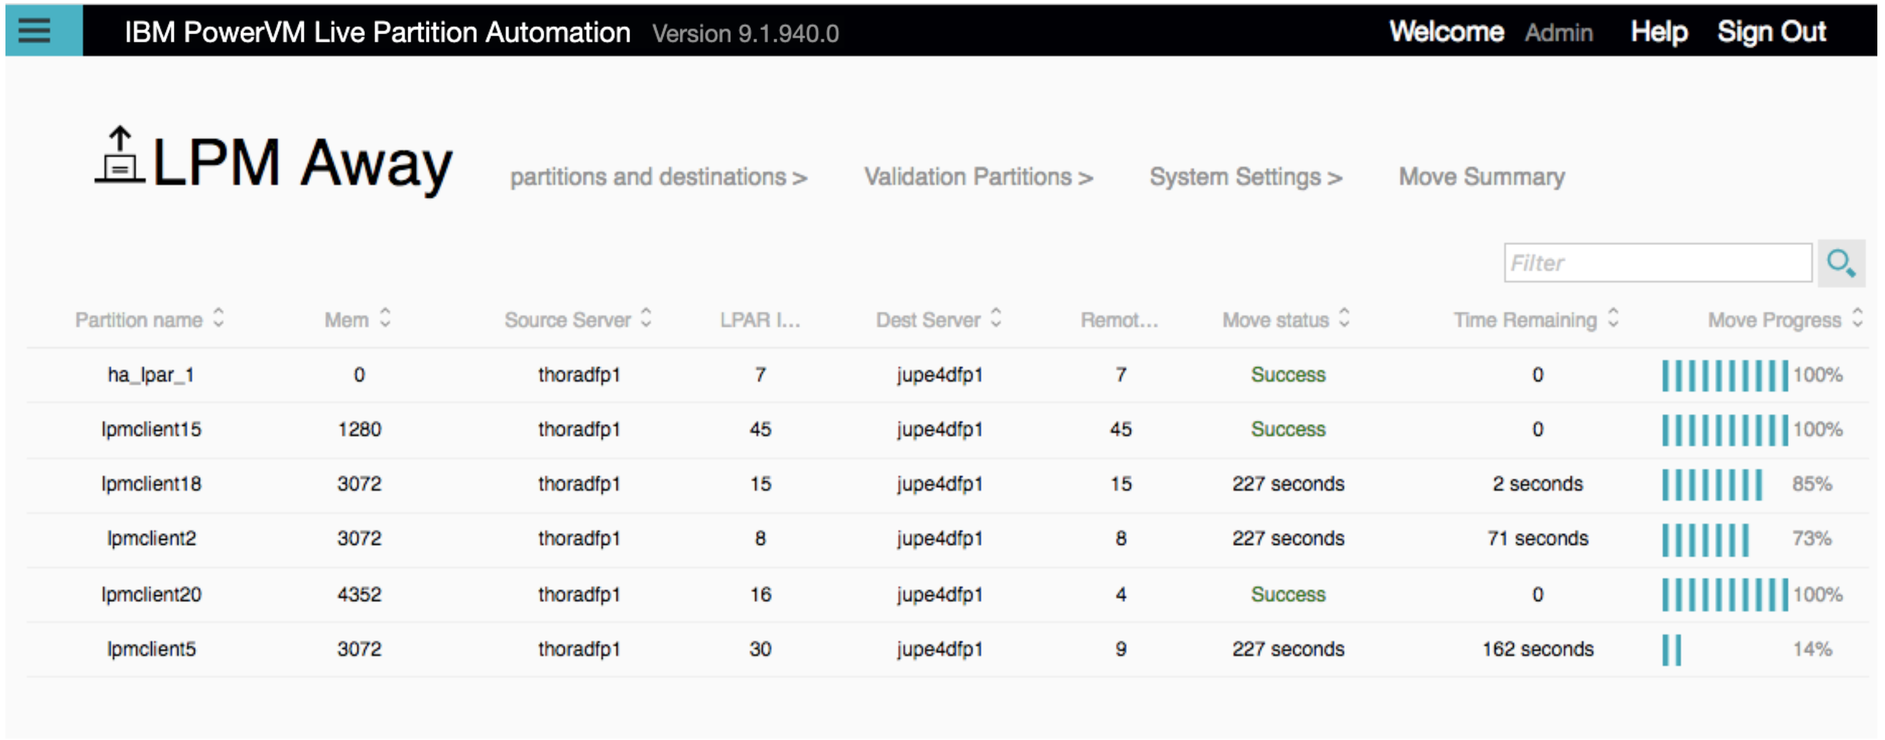

Partition name	Mem	Source Server	LPAR I...	Dest Server	Remot...	Move status	Time Remaining	Move Progress
ha_lpar_1	0	thoradfp1	7	jupe4dfp1	7	Success	0	100%
lpmclient15	1280	thoradfp1	45	jupe4dfp1	45	Success	0	100%
lpmclient18	3072	thoradfp1	15	jupe4dfp1	15	227 seconds	2 seconds	85%
lpmclient2	3072	thoradfp1	8	jupe4dfp1	8	227 seconds	71 seconds	73%
lpmclient20	4352	thoradfp1	16	jupe4dfp1	4	Success	0	100%
lpmclient5	3072	thoradfp1	30	jupe4dfp1	9	227 seconds	162 seconds	14%

Online Resources for the tool

Community website is ibm.biz/lpm_srr_tool

Videos are at ibm.biz/bobtube

PowerVM Provisioning Toolkit v8

PowerVM Provisioning Toolkit v8

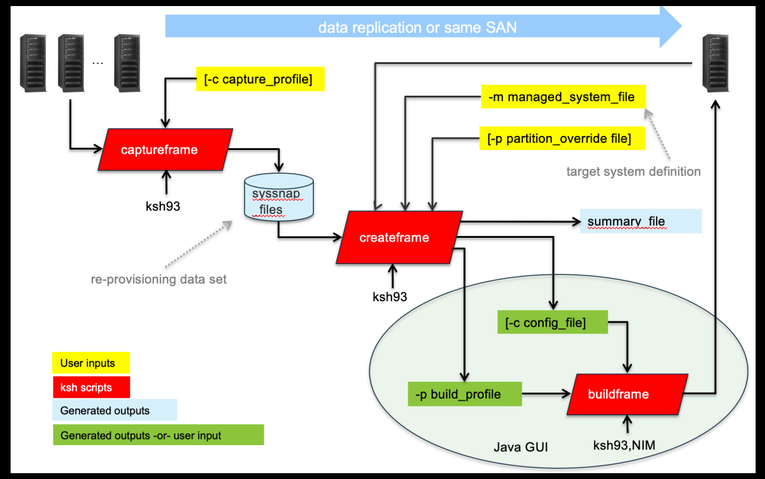

Business Challenge

- Clients struggle with consistency in the build process when deploying VIO servers/clients.
- Differences in configuration depend on who built the systems and when.
- Attempts to develop in-house scripts and automation struggle with adding new features and support

Solution

- The Advanced PowerVM Toolkit to accelerate the client's transition from Power6 → Power7 → Power8 → Power9 across their infrastructure.
- Provide assessment, planning and on-site execution to implement the Toolkit and provisioning process based on client requirements.
- Work with customers to integrate Toolkit into their existing build process and automation

Key Benefits

- Reduced manual errors associated with current build process.
- Create a repeatable methodology for deploying new VIO servers and client LPARs on bare-metal servers.
- Capture entire configuration of an existing system and deploy the same design to other systems. Enabling consistency across datacenters

Manual PowerVM Provisioning Flow

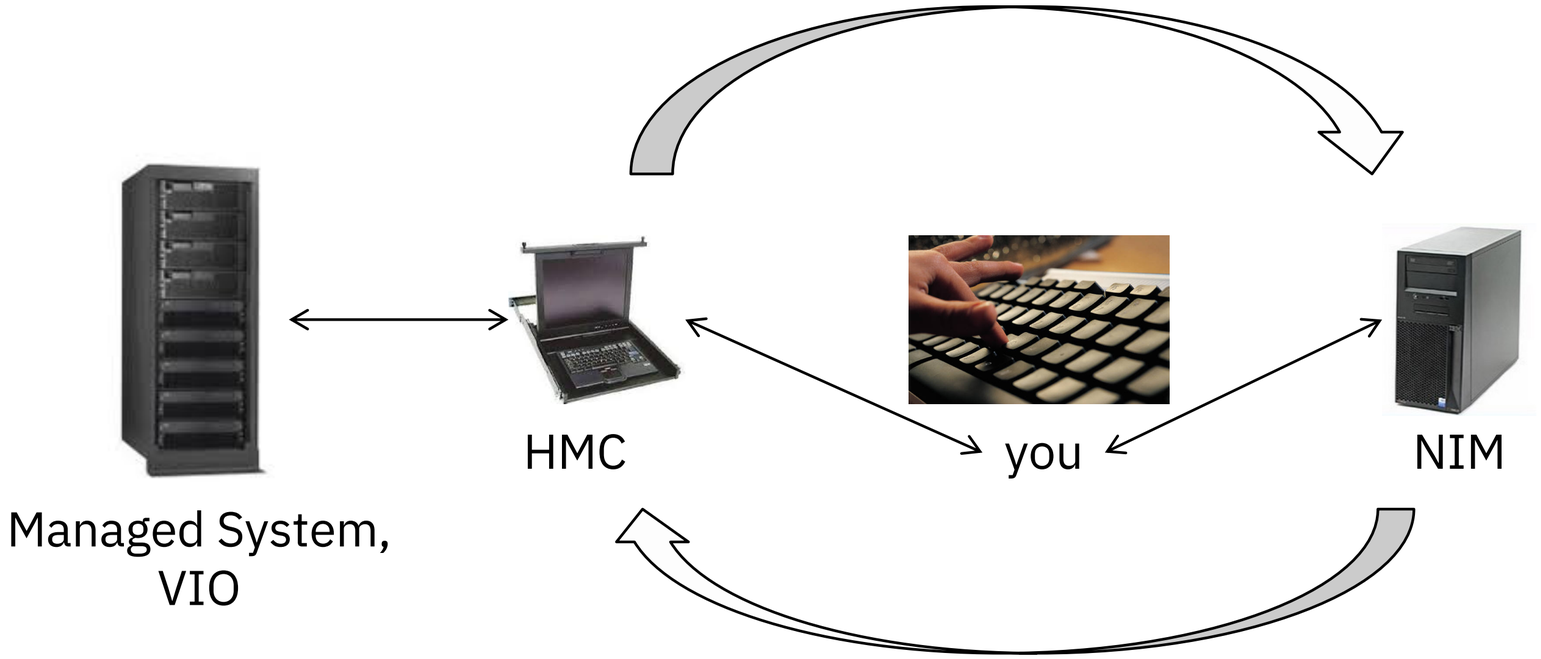

- Gather system build information:
  - HMC, managed system
  - VIO and client names, IP addresses, VLANs, storage, CPU, memory, virtual switches, shared processor pools, SEA/network designs,…
- Repeat a series of ***manual*** interactions using that data:
  - HMC GUI or command line execution
  - NIM commands
  - VIO commands
- ***Consistency can be a struggle…***

IBM Advanced PowerVM Toolkit v8 - Provisioning

buildframe.ksh

- The buildframe script is the original Provisioning Toolkit component dating back to v1 (circa. 2009)
- Provisions VIO servers and clients from easy-to-create ASCII text design patterns (‘*build_profiles*’, ‘*config_files*’, and ‘*include_files*’).
- Speeds and simplifies PowerVM provisioning and gives repeatable, consistent results every time.
- Can be used stand-alone with user-generated, monolithic build profiles.
- With wrapper scripting and config_files/include_files, can be integrated further into any environment for automated, build-on-demand possibilities.

Toolkit ‘GUI’/Editor

- Added to Toolkit in v2 timeframe (since 2012) to assist with build_profie language.
- Editor/syntax assistance for build_profile/config_file/include_file inputs to buildframe.ksh.
- **KEYWORD** samples
- buildframe.ksh execution assistance – work on files and provision right from your laptop!

Toolkit ’wiki’

- ‘TiddlyWiki’ v5 used for Toolkit documentation.
- Embedded graphics
- Searchable/browsable single file - syntax and function reference for all Toolkit files and components.

IBM Advanced PowerVM Toolkit v6 - Provisioning

buildframe.ksh – virtual LPAR builds

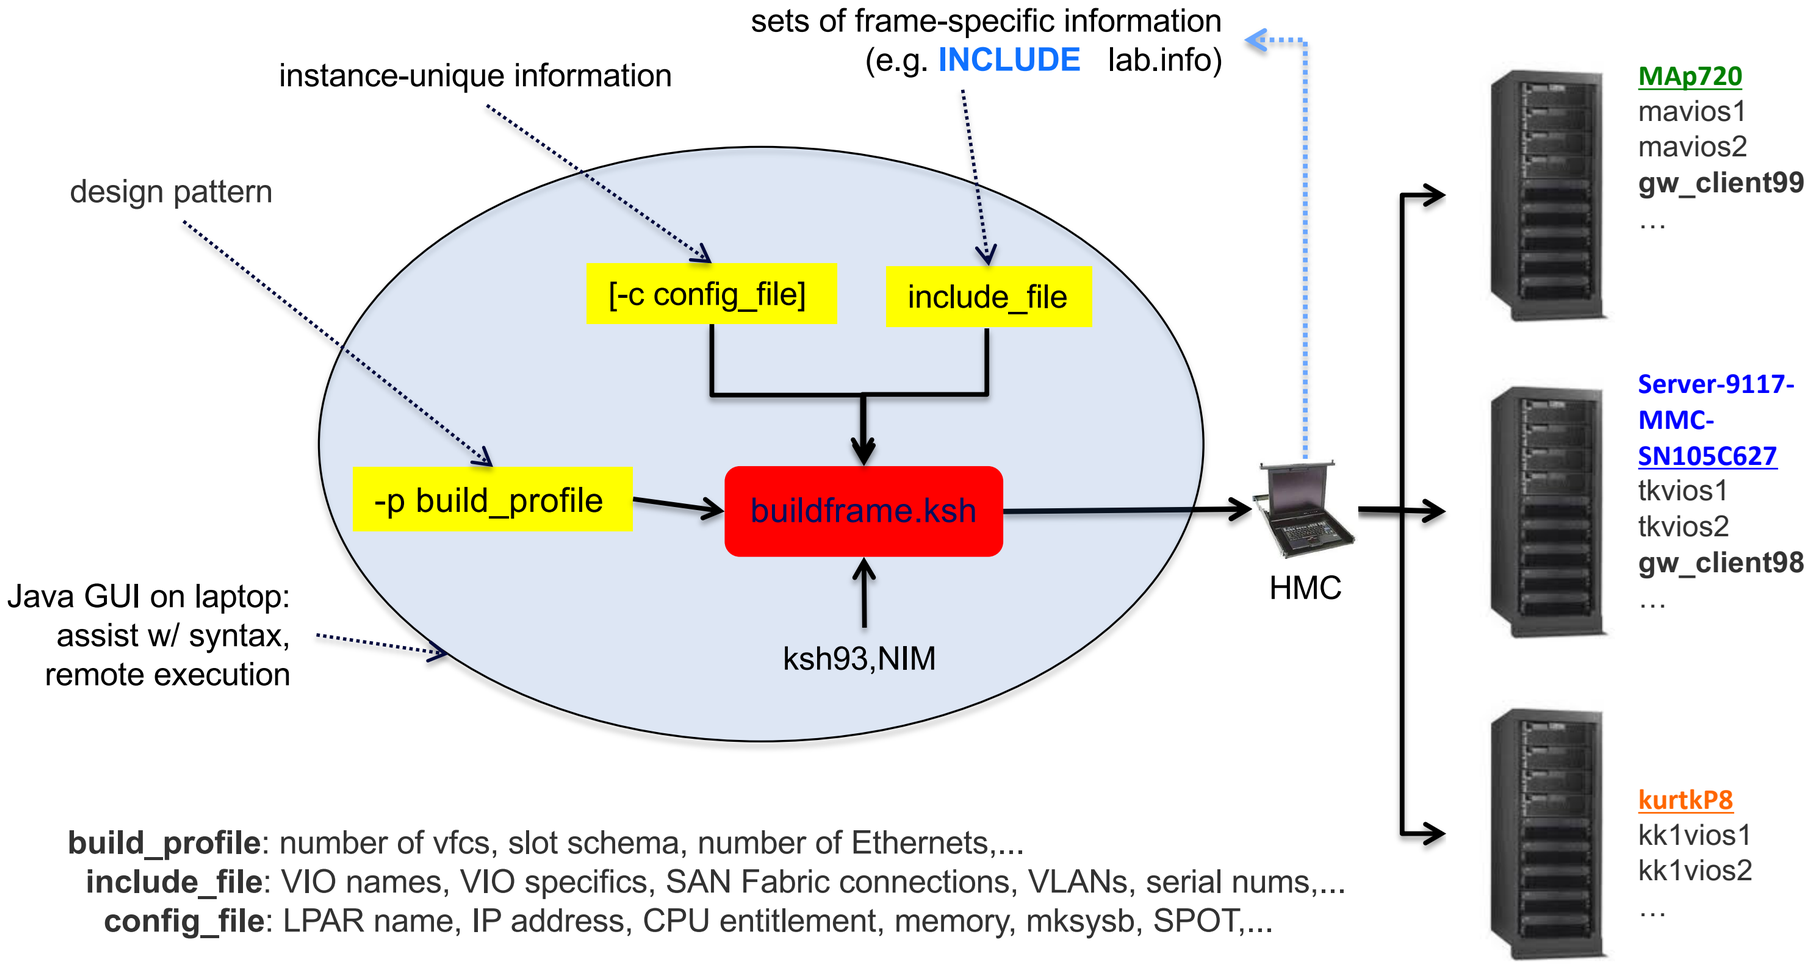

build_profile: number of vfcs, slot schema, number of Ethernets,...
include_file: VIO names, VIO specifics, SAN Fabric connections, VLANs, serial nums,...
config_file: LPAR name, IP address, CPU entitlement, memory, mksysb, SPOT,...

```
./buildframe.ksh  -p lab_LPAR.bp  -i Server-9117-MMC-SN105C627  -f carve,vfcmap,wwpn  -c gw_client98.cfg
./buildframe.ksh  -p lab_LPAR.bp  -i MAp720,VSCSI              -f build              -c gw_client99.cfg
```

buildframe.ksh – Two different modes, -f and -v

In **–f** (**frame** mode), buildframe.ksh creates frame-level objects:

- Partitions (VIOS or VIOC), vswitches, shared processor pools, virtual devices
- ***'-f build'***
  - *validate*
  - *spp*, *vswitch*, *hea*, *carve*
  - *vfcmap*, *wwpn*
  - *vscsimap*, *vscsimapssp*, *mkvg*, *vscsimaplv*
  - *nim*
- 'Utility functions'
  - *nportlogin*, *nportlogout*, *nportquery*
  - *delete*, *shutdown*, *activate*
  - *add*, *netboot*, *drcarve*, *onlycarve*, *onlynim*, *cleannim*

In **-v** (**vio** mode), buildframe.ksh configures VIO servers:

- ***'autoconfig'*** uses same files as used for –f level creation of VIO LPARs.
- Build files added to lpar_netboot payload during `–f nim` – execute automatically on VIO server from /tmp via inittab following mksysb install (can be run manually).
- ***'-v build'***
  - *enattr*, *entattr*, *inet0attr*, *fcsattr*, *fscsiattr* – Set device attributes.
  - *network* – Create Etherchannels, SEAs, assign IP addresses.
  - *mirror* - Initiate VIOS rootvg mirroring.
  - *validate* – Optionally ping to specified gateway(s) from all VIOs IPs.

Use Case 1 – PowerVM VIO Provisioning

Typical Use Case 1 – VIO Server on a net-new frame

- Create VIO server build profile and process with buildframe.ksh:
  - Frame-by-frame specifics can be input via config_files or include_files.
  - Create virtual switches and/or shared processor pools as needed
  - Carve VIO partitions and their virtual devices, assign physical I/O
  - lpar_netboot VIOS LPARs from NIM

  - Optional auto-configuration of VIO includes:
    - Device attributes (en, ent, hba, fcs, fscsi, inet0…)
    - Networks (Etherchannels, Shared Ethernet Adapters, physical Ports)
    - Optional mirror of rootvg
    - pre- and post- build hooks are provided for customer-specific post-build customization

- Upon completion, VIO servers are loaded, configured, running and ready for client LPARs – either for LPM, other migration methods, or new LPAR workloads.

- **All performed with a single command**:

  /buildframe.ksh -p lab_VIOS.bp **-i Server-9117-MMC-SN105C627** -f build

build profile – KEYWORDS

VIO Server build profiles

- **VIO*** (both) # begins a VIO stanza where VIO servers and their owned devices and networks are defined and configured
- SEA (VIO build profiles only)
  - **SEAATTR** # create SEA attribute set (e.g. large_receive=yes)
  - **SEADEF** # create common SEA design used by multiple VIOs
  - **SEA*** # create a Shared Ethernet Adapter from an SEA design
- VIO Physical Networking (VIO only)
  - **PORT*** # Configure IP on a physical *or* virtual ethernet adapter
  - **ETHERCHANNEL*** # create an Etherchannel with optional backup (NIB)
  - **FILERESTORE*** # save/restore TCP/IP files around network config
- Device Attributes (VIO only)
  - **ENATTR** -> **EN*** # apply attributes to en devices
  - **ENTATTR** -> **ENT*** # ent
  - **FCSATTR** -> **FCS*** # fcs
  - **FSCSIATTR** -> **FSCSI*** # fscsi
  - **ECHATTR** -> **ETHERCHANNEL*** # Etherchannel
  - **INET0ATTR** -> **INET0*** # inet0 in ODM
  - **DISKATTR** -> **DISK*** # VIO hdisk, and mirror disk attributes
  - **HBAATTR** -> **HBA*** # configure HBA attributes for VIO
  - **ITM*** # configure ITM VIO VA premium agent

Toolkit UI/build profile editor

build_profile/config_file syntax assistance

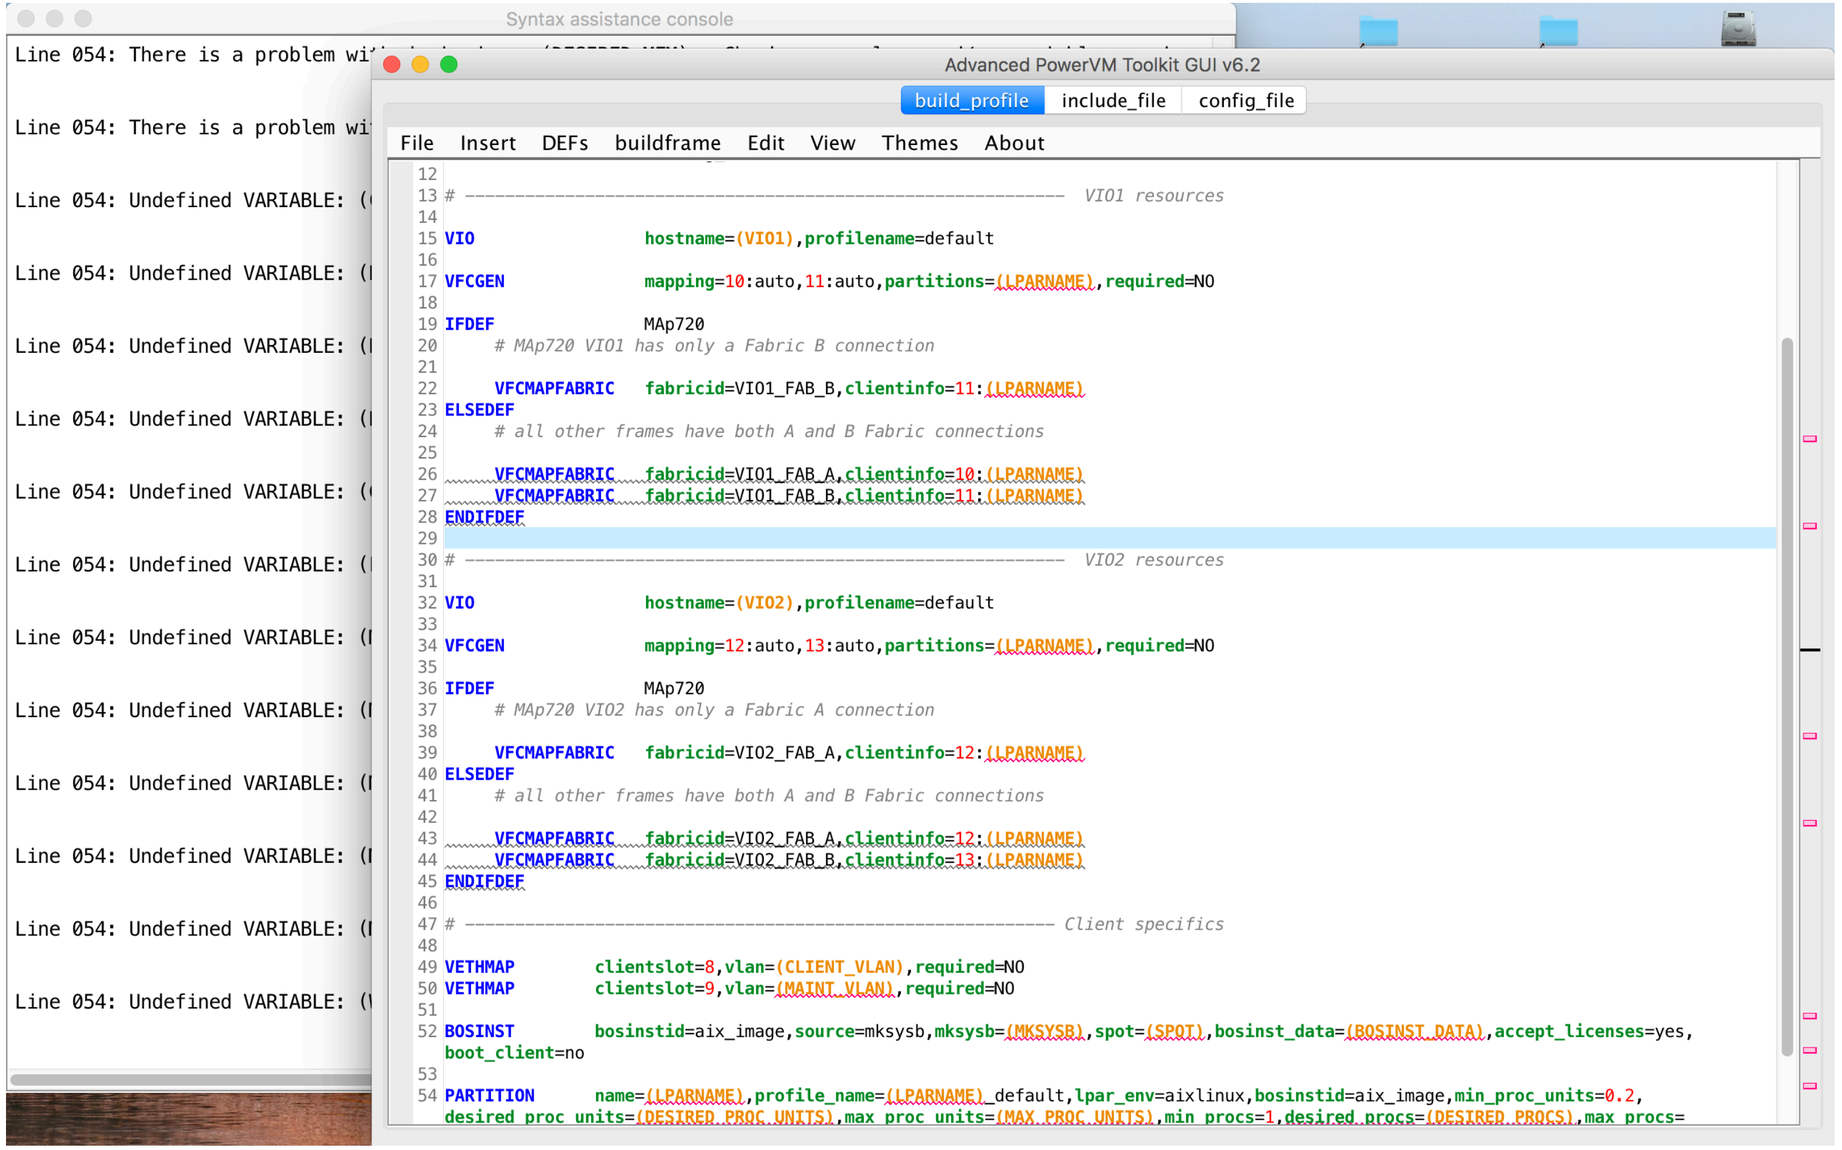

Use Case 2 – PowerVM Client LPAR Provisioning

Typical Use Case 2 – VIO Client provisioning

- VIO servers/frame need not have been built with the Toolkit in order to use it to deploy new LPARs.
- Create client LPAR build profile for one or N client partitions and process with buildframe.ksh to:
  - Carve client LPARs and virtual SCSI, NPIV, Ethernet devices, VNIC
  - Map VSCSI and/or NPIV storage (even Shared Storage Pools)
  - Optionally assign physical I/O (not recommended)
  - lpar_netboot AIX client LPARs from NIM
- As with VIO servers, frame specifics can be isolated into **include_files**, with LPAR specifics represented in **config_files**.

build_profile sample – KEYWORDS with key=values

SERVERINFO servername=Server-9117-MMC-SN105C627,buildhmc=mghmc.rchland.ibm.com

VETHMAP clientslot=2,vlan=50

FABRIC name=FabricA,clientadapters=4
FABRIC name=FabricB,clientadapters=7

VIO hostname=tkvios1,profilename=normal

VFCGEN mapping=4:auto
VFCMAP port=DBJF101-P2-C1-T1,clientinfo=4:tk_client1

VIO hostname=tkvios2,profilename=normal

VFCGEN mapping=7:auto
VFCMAP port=DBJF103-P2-C1-T4,clientinfo=7:tk_client1

BOSINST bosinstid=aix71,source=mksysb,mksysb=mksysb_AIX710TL02SP2,
spot=AIX710TL2SP2spot, bosinst_data=bosinst_data_aix,accept_licenses=yes,boot_client=no

PARTITION name=tk_client1,profile_name=normal,min_mem=6144,desired_mem=16384,
max_mem=24576,min_proc_units=0.1,desired_proc_units=0.3,max_proc_units=2.0,min_procs=1,
desired_procs=2,max_procs=3,lpar_env=aixlinux,proc_mode=shared,uncap_weight=128,mem_mode=ded,
allow_perf_collection=1,sharing_mode=uncap,max_virtual_slots=11,all_resources=0,conn_monitoring=1,
redundant_err_path_reporting=0,**bosinstid**=aix71,
lpar_netboot=10.0.0.11,tk_client1,10.0.0.55,255.255.255.0,10.0.0.1,NA,auto,auto,NA

buildframe.ksh 101

NPIV Client LPAR creation

Assuming the build profile on the previous page is named *tk_client1.def*, we might build that NPIV client this way…

1. buildframe.ksh –p *tk_client1.def* **-f carve** # create partition profile and devices
2. buildframe.ksh –p *tk_client1.def* **-f vfcmap** # map vfcs to VIOS-owned fibre ports
3. buildframe.ksh –p *tk_client1.def* **-f wwpn** # ‘harvest’ WWPNs for SAN team
4. …Obtain SAN zoning and allocation…
5. buildframe.ksh –p *tk_client1.def* **-f nim** # lpar_netboot from NIM

Or… buildframe.ksh –p *tk_client1.def* **-f carve,vfcmap,wwpn**
and then
buildframe.ksh –p *tk_client1.def* **-f nim**

Use Case 3 – ‘Re-Provisioning’

IBM Advanced PowerVM Provisioning Toolkit v3+

captureframe/createframe/buildframe

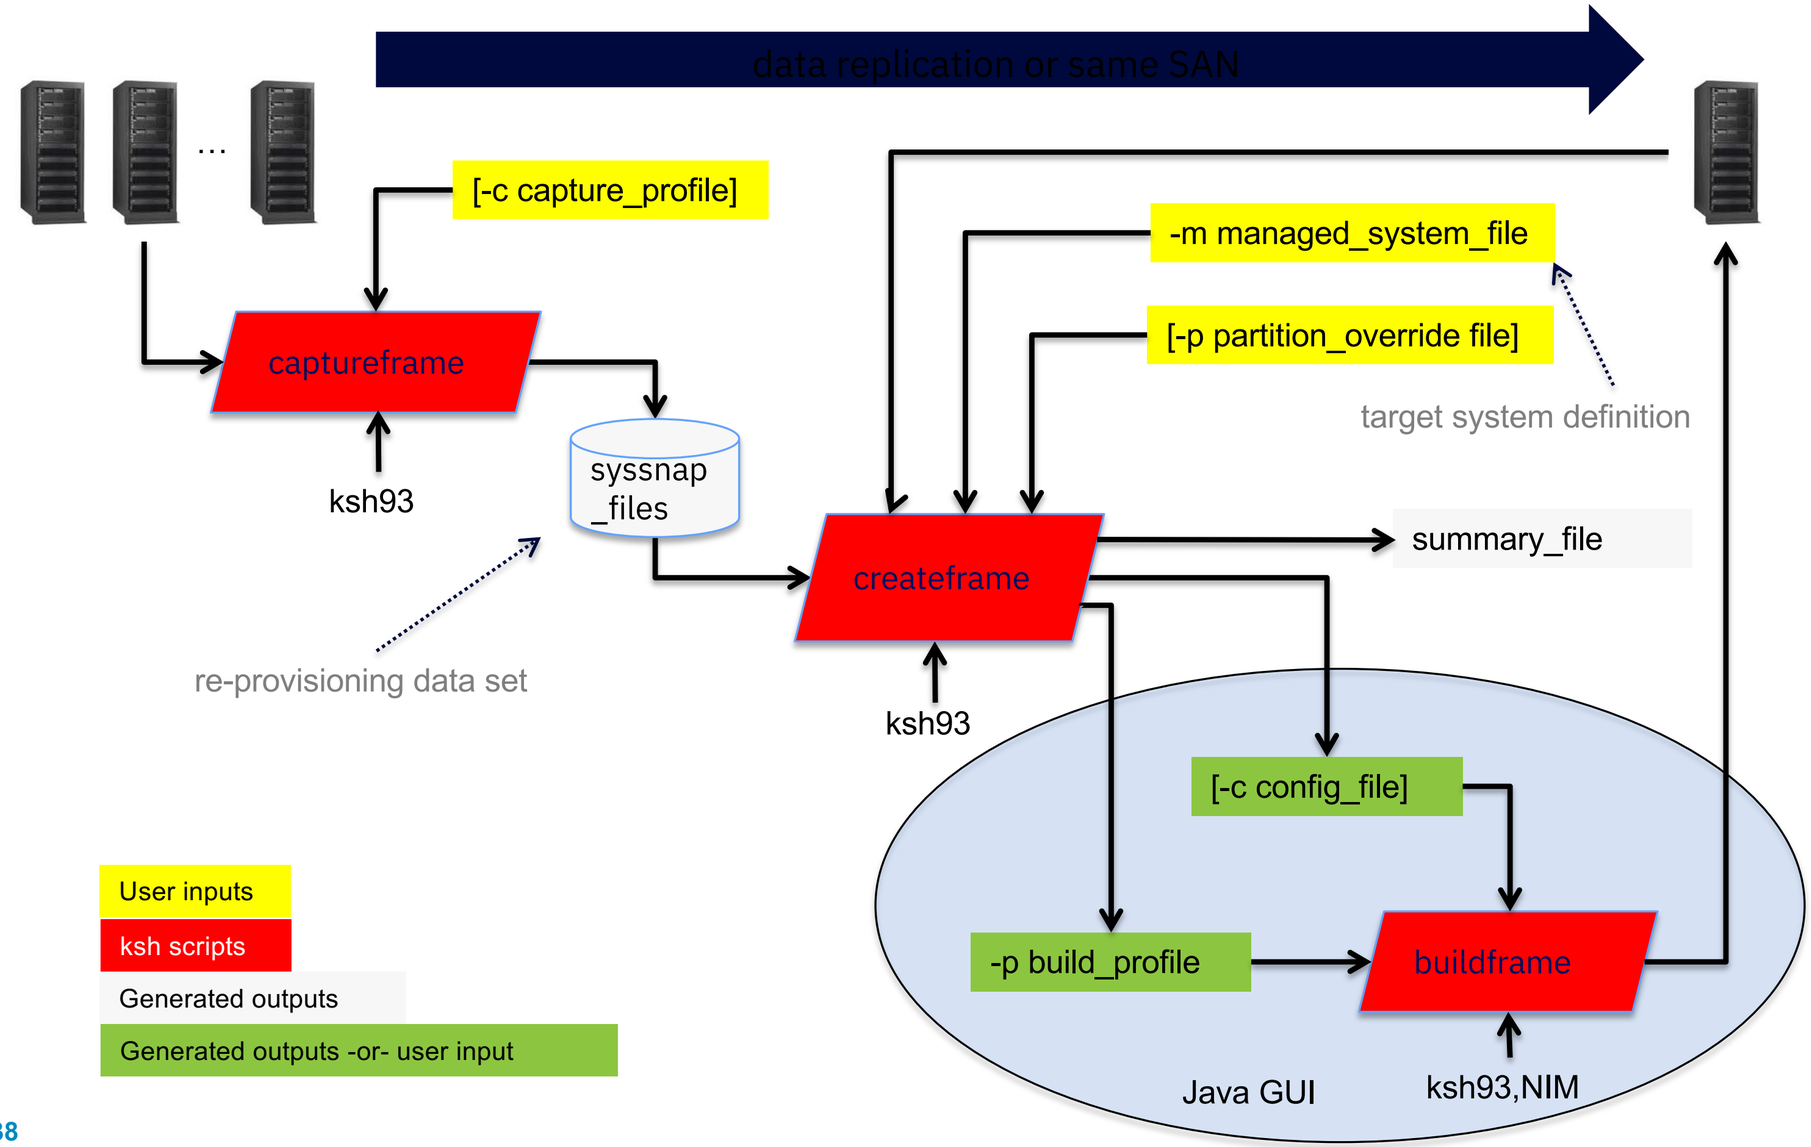

38

Capture/create/build application to LPAR migration

Basic use case

1. Capture a source group of LPARs to move to a common target frame with ***captureframe.ksh***.
2. Validate captured data – no files should be missing or contain errs.
3. Prepare a managedsystem_file to define a target frame and its mapping scheme.
4. Prepare partition_file to override any specific LPAR properties, if needed.
5. Process captured data, managedsystem_file, partition_file through ***createframe.ksh***.
6. Examine summary file and resulting config_file and build_profile.
7. Adjust build_profile(s) and config_file as required or change managedsystem_file and return to step 5.
8. Process client build profile and config file through ***buildframe.ksh*** to stage new LPARs.
9. Validate new LPAR shells on target frame. Delete and restage if needed.
10. Deactivate source LPARs and activate target LPARs as per plan. (“toggle”)

The Provisioning Toolkit does not migrate data; it ‘re-provisions’ LPARs:

– It wires up existing storage known *a priori* (same or replicated for VSCSI),
– Duplicates virtual fibre channel WWPN values (virtual MAC address, too)

– Can also builds all-new LPAR shells for mksysb-style installation/migration

Use Case 4 – LPAR de-commission

IBM Advanced PowerVM Provisioning Toolkit v5+

- ***delete_lpars.ksh***
  - Simplifying wrapper script around the powerful `buildframe –f delete` function.
  - May be used whether or not the Toolkit was used to build the LPAR in the first place.
    - Dynamically generates build_profile for input to buildframe.
  - Same functional environment and directories as for Provisioning Toolkit core scripts.
  - Uses same user(s) and SSH key(s).
  - Assumes short, symbolic names to existing Toolkit scripts.
    - buildframe.ksh -> buildframe.ksh.v6.0
  - Simple calling convention: single HMC, single frame, prompt or no-prompt, single LPAR name
    - Optional HMC user (default hscroot).
    - Checks/requires that LPAR is ‘Not Activated’
    - Unmaps and deletes any vfchost, vhost devices on parent VIO servers then saves VIO profiles if needed.
    - Deletes LPAR and its profile.

./delete_lpar.ksh -u toolkit -h mghmc -m Server-9117-MMC-SN105C627 -n gw_client77

Capture/create/build application to LPAR migration

Basic use case

1. Capture a source group of LPARs to move to a common target frame with ***captureframe.ksh***.
2. Validate captured data – no files should be missing or contain errs.
3. Prepare a managedsystem_file to define a target frame and its mapping scheme.
4. Prepare partition_file to override any specific LPAR properties, if needed.
5. Process captured data, managedsystem_file, partition_file through ***createframe.ksh***.
6. Examine summary file and resulting config_file and build_profile.
7. Adjust build_profile(s) and config_file as required or change managedsystem_file and return to step 5.
8. Process client build profile and config file through ***buildframe.ksh*** to stage new LPARs.
9. Validate new LPAR shells on target frame. Delete and restage if needed.
10. Deactivate source LPARs and activate target LPARs as per plan. (“toggle”)

The Provisioning Toolkit does not migrate data; it ‘re-provisions’ LPARs:

– It wires up existing storage known *a priori* (same or replicated for VSCSI),
– Duplicates virtual fibre channel WWPN values (virtual MAC address, too)

– Can also builds all-new LPAR shells for mksysb-style installation/migration

SR-IOV Migration Tool

SR-IOV Migration Tool

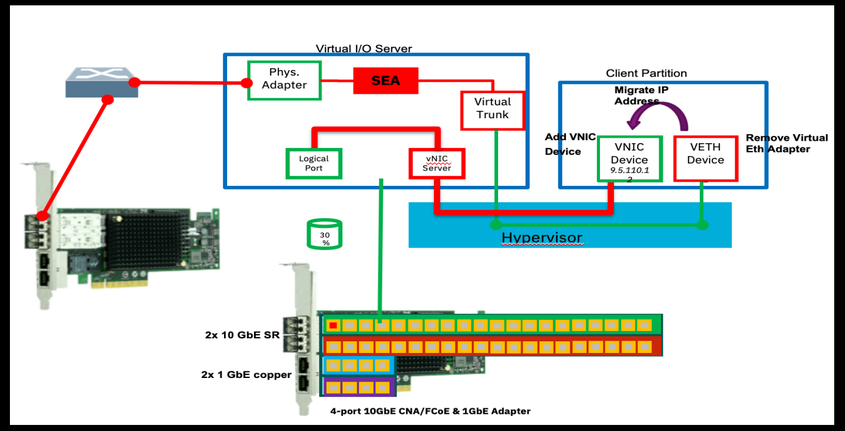

Dynamically add a new SR-IOV VNIC adapter to a running client LPAR

- Currently supports AIX versions 7.1 & 7.2
- Linux support planned for future release, starting with SUSE Linux on Power

Live-migrate IP address from a specified virtual Ethernet adapter to the new VNIC

- Non-disruptive IP address migration while the LPAR is up and running (AIX only)
- Linux IP address migration planned (disruptive)

Remove existing virtual Ethernet adapter

- Optionally (dynamically) remove virtual Ethernet adapter.
- Removes OS level-devices, as well as saves configuration to client's HMC partition profile

Traditional Network Options for POWER Systems

- **Dedicated Adapters**
  - Best possible performance
  - Adapter exclusively bound to particular partition; no resource sharing
- **Virtual Ethernet Adapter**
  - Hypervisor internal switching
- **VIOS Shared Ethernet Adapter**
  - Hypervisor Switch “uplink” to physical network through Virtual I/O Server (required).
  - Options for high availability
    - SEA failover, SEA failover w. load sharing, Network Interface Backup (NIB)
- **Why do we need something new?**
  - Overcome software based virtualization limitations.
  - Adoption of high network bandwidths of 10 Gigabit Ethernet and beyond.
  - Running high network speeds for affordable overhead.

PowerVM SR-IOV VNIC Overview

Single Root I/O Virtualization (SR-IOV) *defines extensions to the PCI Express® specification to allow multiple partitions to share a PCIe® device.*

vNIC *is a new virtual adapter type, based on SR-IOV technology.*

Values:

- *Increased I/O efficiency*
- *Better “out of the box” performance*
- *Maximize use of high bandwidth I/O infrastructure and adapters*
- *Facilitates consolidation by reducing data center hardware requirement, associated energy costs, and floor space requirements*

SR-IOV Migration Tool

AIX client LPAR IP address migration

Moving IP Addresses seamlessly with AIX `ifconfig transfer`

```
# netstat -in
Name  Mtu   Network     Address              Ipkts Ierrs    Opkts Oerrs  Coll
en0   1500  link#2      6e.14.ca.5b.9.4    8275945      0  3937496      0     0
en0   1500  9.156.175   9.156.175.25       8275945      0  3937496      0     0
lo0   16896 link#1                         21583      0    21583      0     0
lo0   16896 127         127.0.0.1          21583      0    21583      0     0
lo0   16896 ::1%1                         21583      0    21583      0     0
```

```
# ifconfig en0 9.156.175.25 transfer en1
```

```
# netstat -in
Name  Mtu   Network     Address              Ipkts Ierrs    Opkts Oerrs  Coll
en0   1500  link#2      6e.14.ca.5b.9.4    8276617      0  3937572      0     0
en1   1500  link#3      c6.e6.3.d7.79.0         16      0        3      0     0
en1   1500  192.168.11  9.156.175.25            16      0        3      0     0
lo0   16896 link#1                         21628      0    21628      0     0
lo0   16896 127         127.0.0.1          21628      0    21628      0     0
lo0   16896 ::1%1                         21628      0    21628      0     0
```

SR-IOV Migration Tool

Live-migration from virtual Ethernet to SR-IOV devices

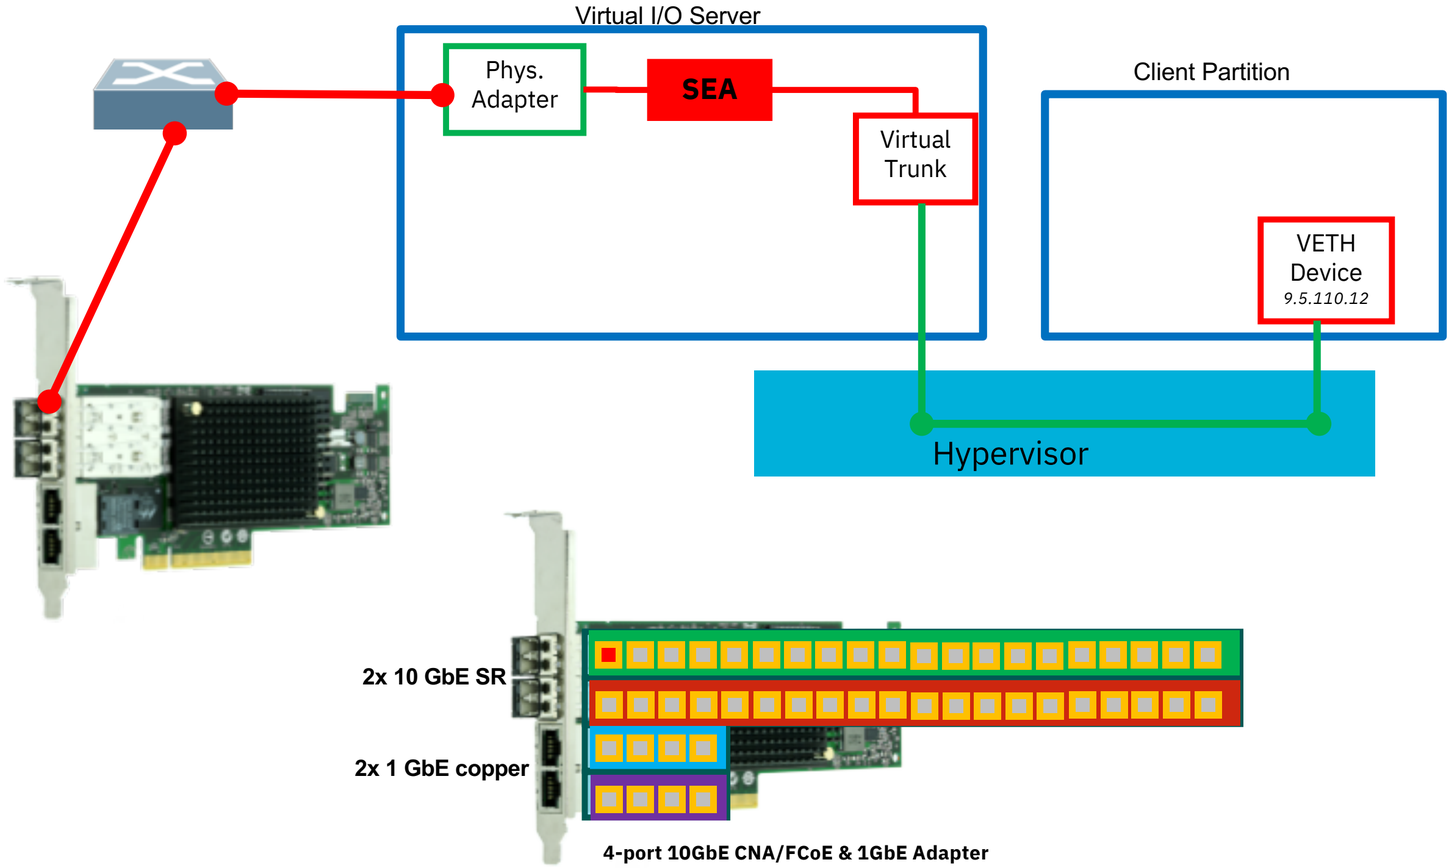

SR-IOV Migration Tool

Live-migration from virtual Ethernet to SR-IOV devices

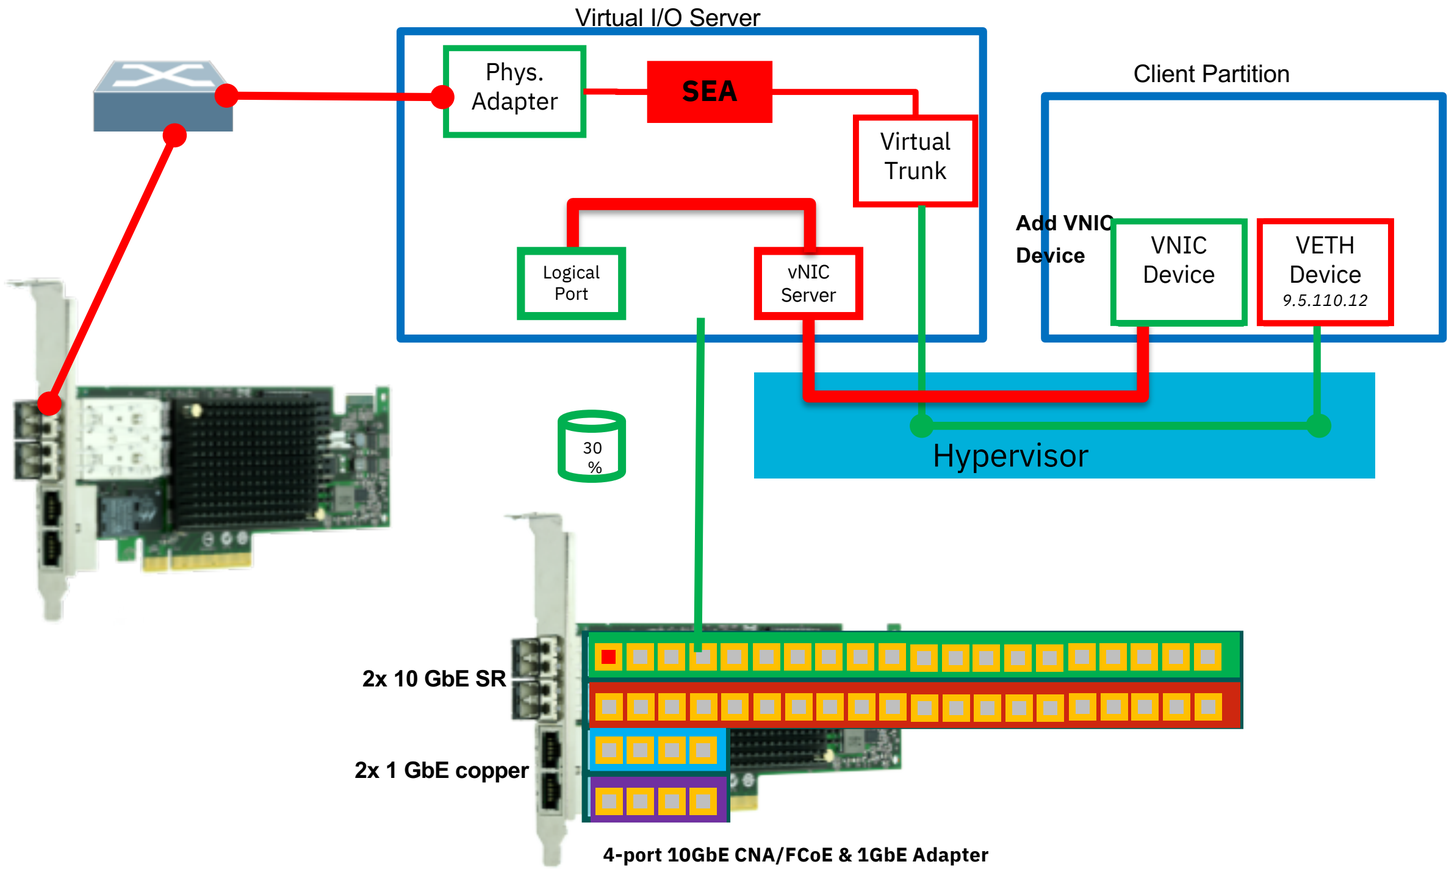

SR-IOV Migration Tool

Live-migration from virtual Ethernet to SR-IOV devices

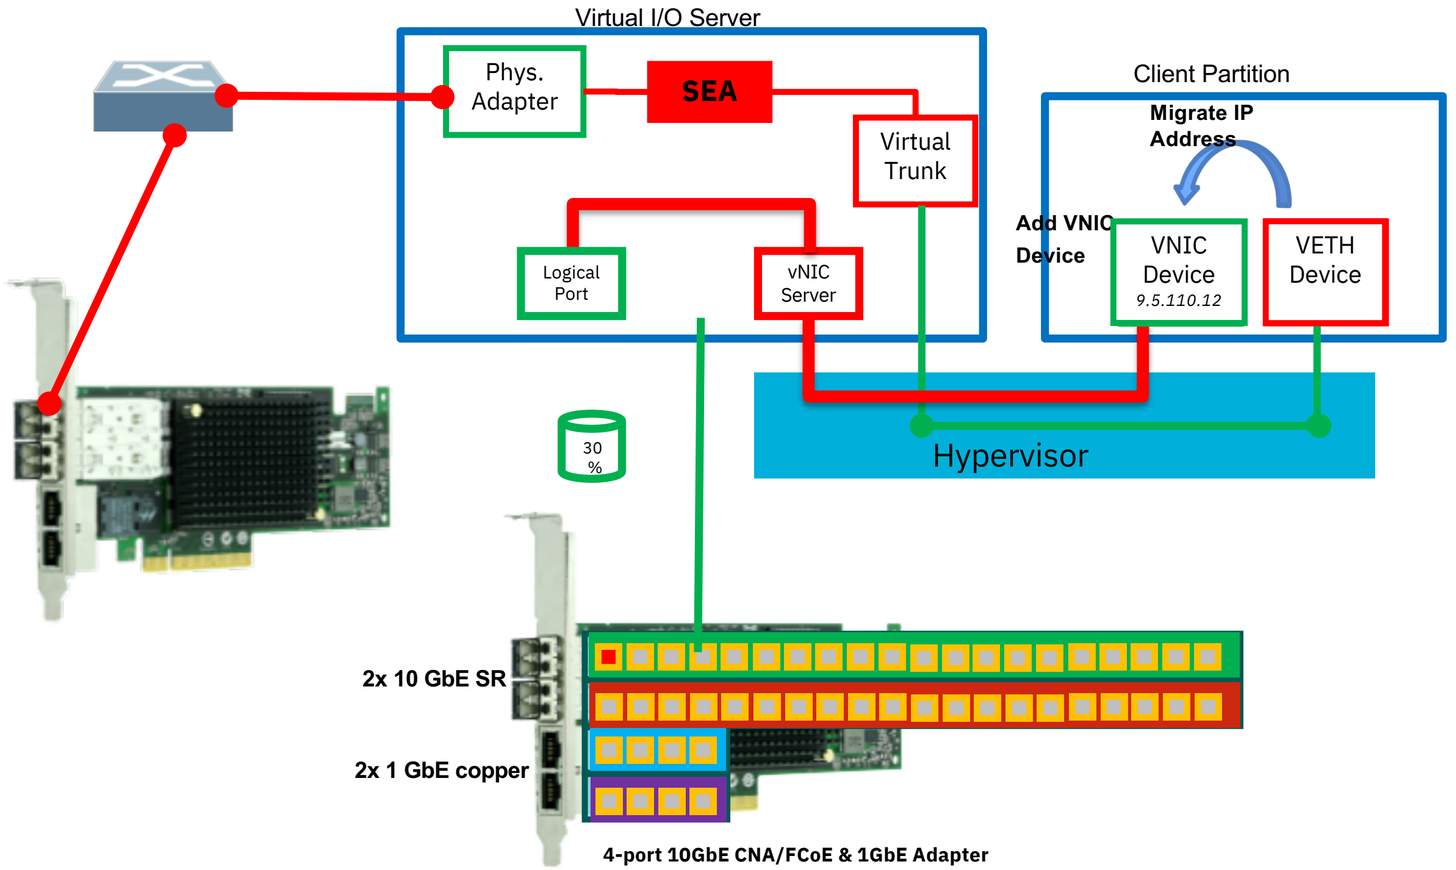

SR-IOV Migration Tool

Live-migration from virtual Ethernet to SR-IOV devices

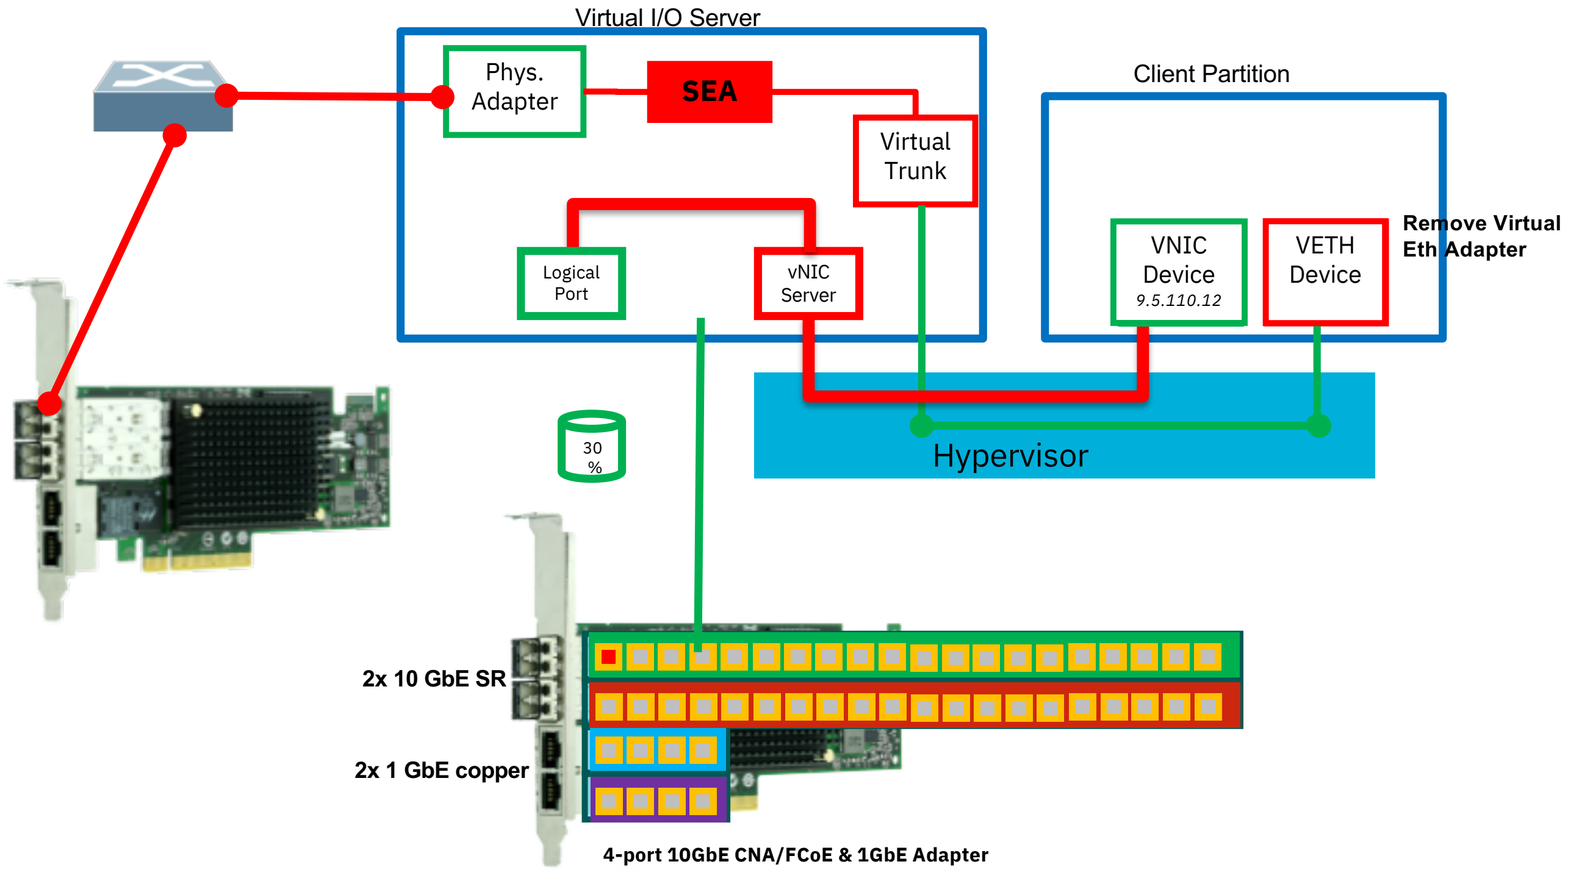

SR-IOV Migration Tool

Live-migration from virtual Ethernet to SR-IOV devices

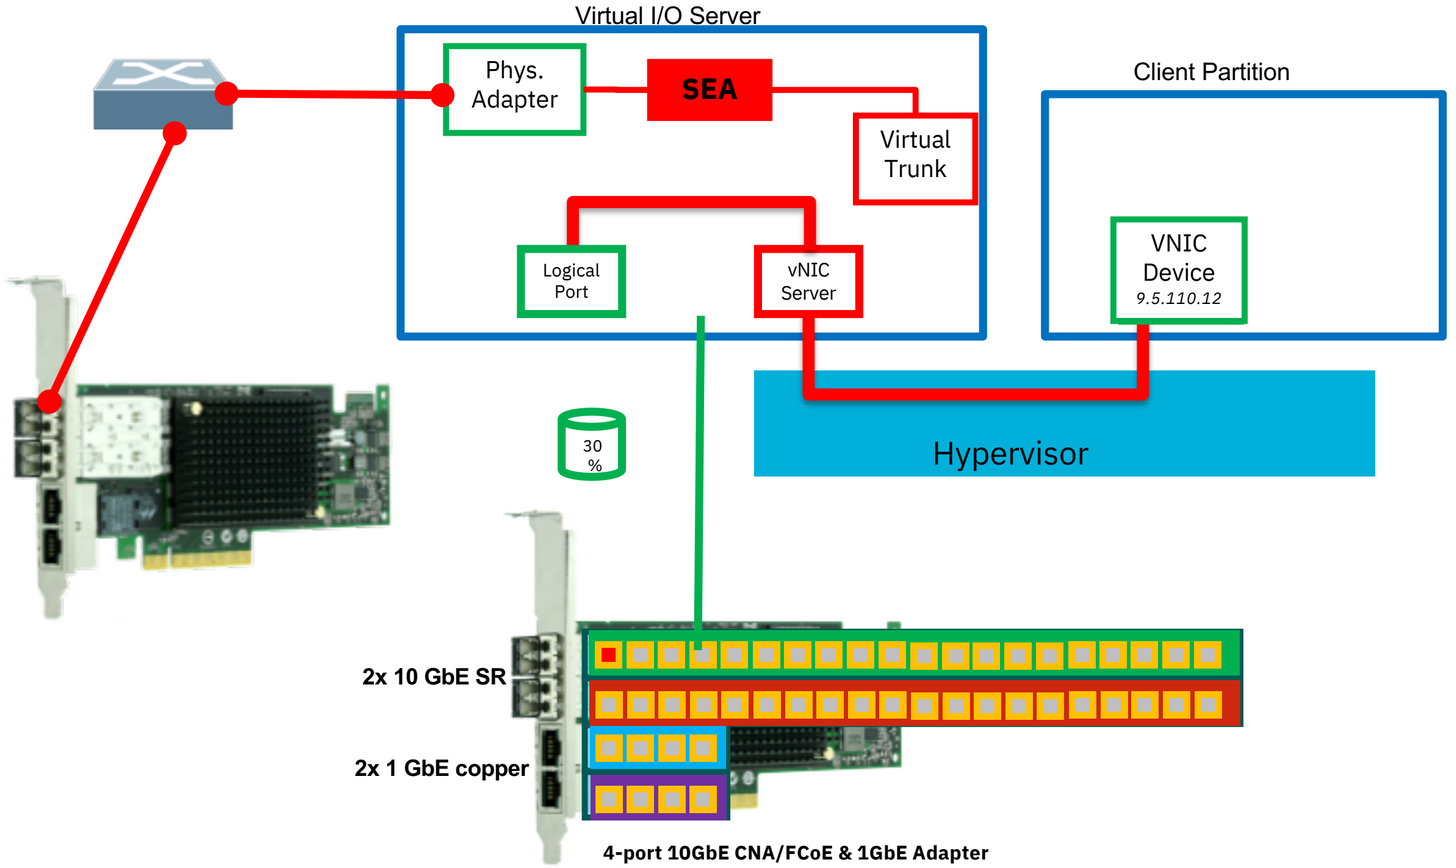

VSCSI-to-NPIV Migration Tool

VSCSI-to-NPIV Migration Tool

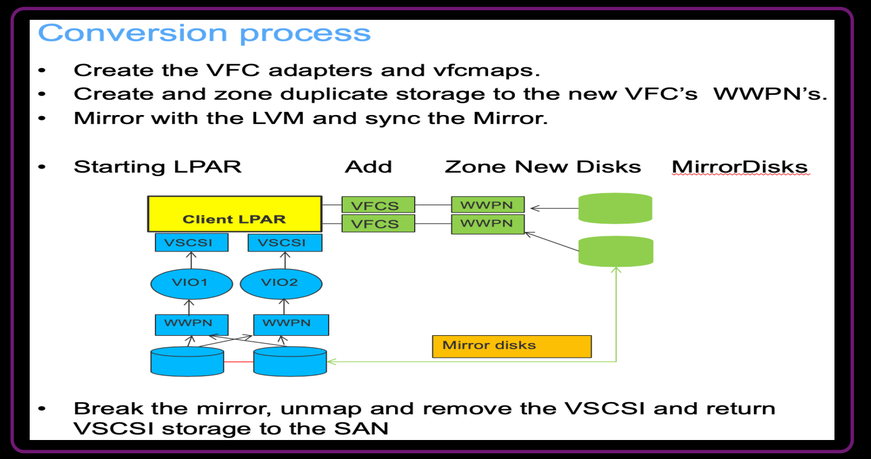

Convert a VSCSI partition to NPIV with **no downtime**

NPIV virtual adapter creation/mapping

LUN mirroring and mirror breaking

VSCSI virtual adapter unmapping/deletion

VSCSI-to-NPIV Migration Tool

Accelerate conversion of a VSCSI partition to NPIV with **no downtime**

- NPIV virtual adapter creation/mapping
- LUN mirroring and mirror breaking
- VSCSI virtual adapter unmapping/deletion
- Support for EMC, 3PAR, NAS, Hitachi and IBM storage
- PowerHASupport

VSCSI to NPIV Menu

1: Get LUNid's and LUN Sizes
2: Create VFC Adapters
3: Collect WWPN's
4: Create VFC Mappings
5: Login/Logout your WWPNs
6: Set Disk Attributes
7: MirrorVG
8: Break the Mirror
9: Clean up VSCSI Mappings
10: PowerHA Options
11: Exit

Conversion process

- Create the VFC adapters and vfcmaps.
- Create and zone duplicate storage to the new VFC's WWPN's.
- Mirror with the LVM and sync the Mirror.

- Starting LPAR Add Zone New Disks MirrorDisks

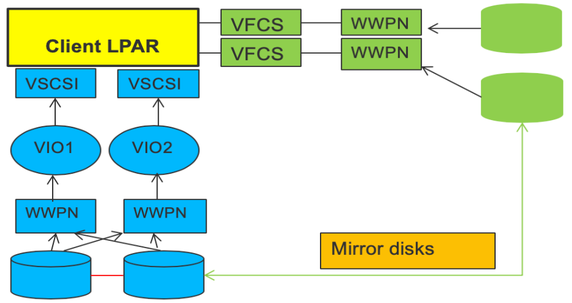

- Break the mirror, unmap and remove the VSCSI and return VSCSI storage to the SAN

Capacity on Demand Dashboard (CoDD)

Capacity on Demand Dashboard

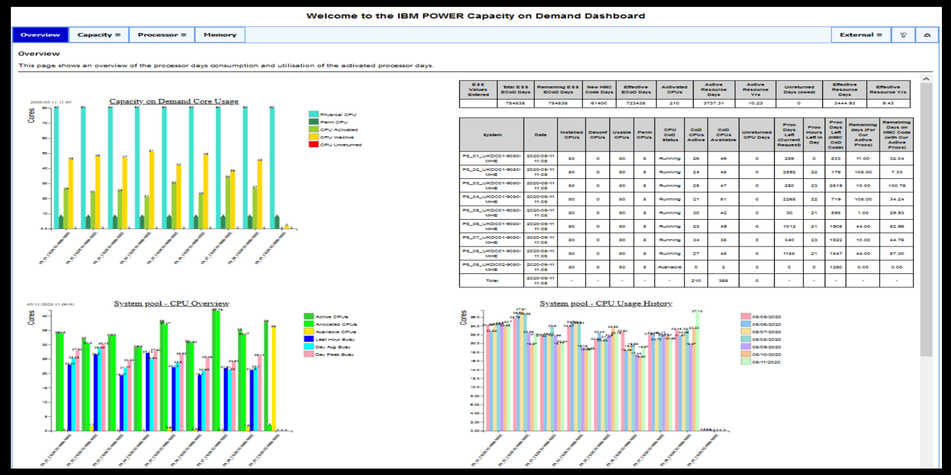

Consolidates data from multiple HMCs into a single dashboard.

- Simplifies access to HMC data and provides context to help analyse Elastic CoD activations

Capacity on Demand Status

- Snapshot of status for each server showing where applied codes require refreshing.
- Breakdown of server activation codes with consumption rate at current resource level.
- Consolidates server data to indicate the consumption of pre-purchased days

Utilisation Overview

- System Overview to compare resource assignment against actual & peak usage
- Historical Utilisation charts show trends in resource usage

CoD Dashboard - Default View

- The initial display provides an overview of the common Processor information charts.
- Menu options for specific resource views show larger charts and additional data

Displays the active and available CoD resources for each server

Any CPU days owed against the current HMC activation code are flagged by red bars

Remaining CPU days from the Entitled Systems Support site and calculates the burn rate based on the current CPU activations

Individual server burn rate for currently applied & available HMC codes

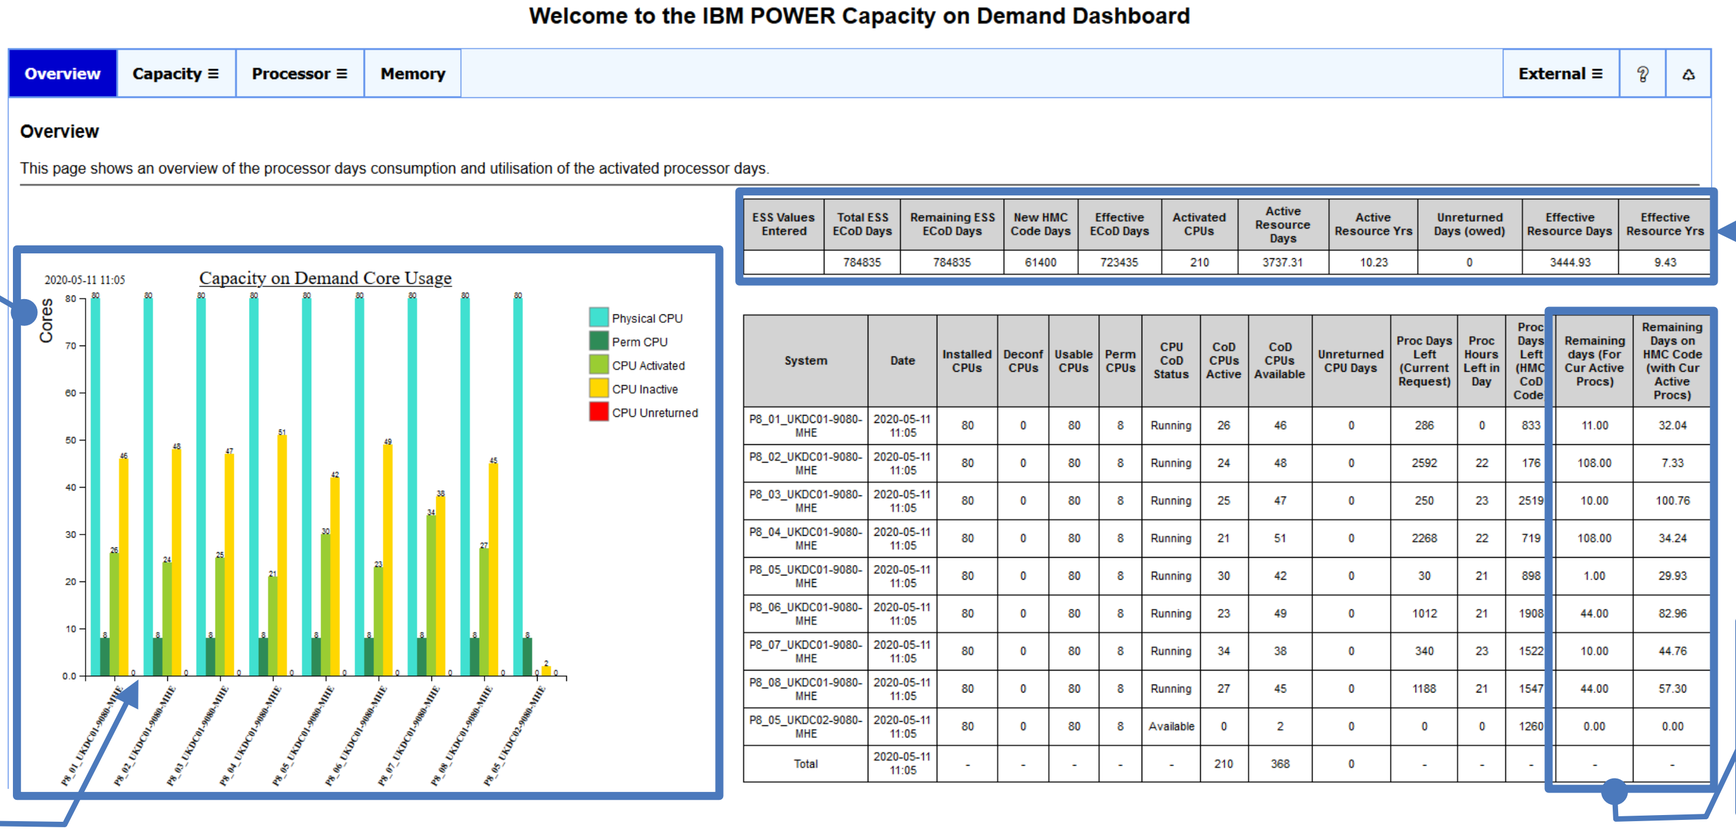

ESS Values Entered	Total ESS ECoD Days	Remaining ESS ECoD Days	New HMC Code Days	Effective ECoD Days	Activated CPUs	Active Resource Days	Active Resource Yrs	Unreturned Days (owed)	Effective Resource Days	Effective Resource Yrs
	784835	784835	61400	723435	210	3737.31	10.23	0	3444.93	9.43

System	Date	Installed CPUs	Deconf CPUs	Usable CPUs	Perm CPUs	CPU CoD Status	CoD CPUs Active	CoD CPUs Available	Unreturned CPU Days	Proc Days Left (Current Request)	Proc Hours Left in Day	Proc Days Left (HMC CoD Code	Remaining days (For Cur Active Procs)	Remaining Days on HMC Code (with Cur Active Procs)
P8_01_UKDC01-9080-MHE	2020-05-11 11:05	80	0	80	8	Running	26	46	0	286	0	833	11.00	32.04
P8_02_UKDC01-9080-MHE	2020-05-11 11:05	80	0	80	8	Running	24	48	0	2592	22	176	108.00	7.33
P8_03_UKDC01-9080-MHE	2020-05-11 11:05	80	0	80	8	Running	25	47	0	250	23	2519	10.00	100.76
P8_04_UKDC01-9080-MHE	2020-05-11 11:05	80	0	80	8	Running	21	51	0	2268	22	719	108.00	34.24
P8_05_UKDC01-9080-MHE	2020-05-11 11:05	80	0	80	8	Running	30	42	0	30	21	898	1.00	29.93
P8_06_UKDC01-9080-MHE	2020-05-11 11:05	80	0	80	8	Running	23	49	0	1012	21	1908	44.00	82.96
P8_07_UKDC01-9080-MHE	2020-05-11 11:05	80	0	80	8	Running	34	38	0	340	23	1522	10.00	44.76
P8_08_UKDC01-9080-MHE	2020-05-11 11:05	80	0	80	8	Running	27	45	0	1188	21	1547	44.00	57.30
P8_05_UKDC02-9080-MHE	2020-05-11 11:05	80	0	80	8	Available	0	2	0	0	0	1260	0.00	0.00
Total	2020-05-11 11:05	-	-	-	-	-	210	368	0	-	-	-	-	-

CoD Dashboard - ‘Capacity’ Views

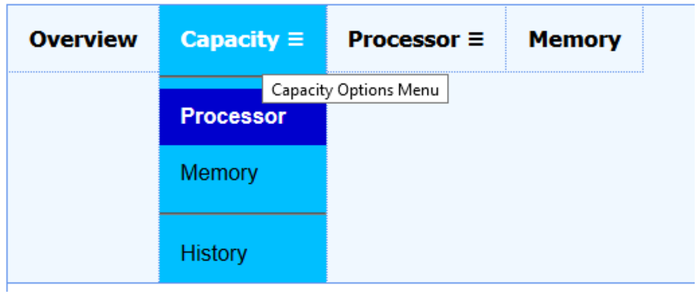

- The **History** view shows the CoD activity from all systems, including what resources were activated by codes applied on the HMC against individual systems.

- The **Processor** & **Memory** views gives the CoD status and the “burn rate” tables as well as additional Capacity charts.

CPU Capacity on Demand Burn

CoD Days manually entered using information from the IBM ESS site
("My entitled hardware" ==> "Elastic CoD - Generate new codes")

ESS Values Entered	Total ESS ECoD Days	Remaining ESS ECoD Days	New HMC Code Days	Effective ECoD Days	Activated CPUs	Active Resource Days	Active Resource Yrs	Unreturned Days (owed)	Effective Resource Days	Effective Resource Yrs
	784835	784835	61400	723435	210	3737.31	10.23	0	3444.93	9.43

Change the CoD days values by running : **~CoDD/config/CoDD_days.sh**

CPU Capacity on Demand Usage per System

Based on information collected from the HMC

System	Date	Installed CPUs	Deconf CPUs	Usable CPUs	Perm CPUs	CPU CoD Status	CoD CPUs Active	CoD CPUs Available	Unreturned CPU Days	Proc Days Left (Current Request)	Proc Hours Left in Day	Proc Days Left (HMC CoD Code)	Remaining days (For Cur Active Procs)	Remaining Days on HMC Code (with Cur Active Procs)
P8_01_UKDC01-9080-MHE	2020-05-11 11:05	80	0	80	8	Running	26	46	0	286	0	833	11.00	32.04
P8_02_UKDC01-9080-MHE	2020-05-11 11:05	80	0	80	8	Running	24	48	0	2592	22	176	108.00	7.33
P8_03_UKDC01-9080-MHE	2020-05-11 11:05	80	0	80	8	Running	25	47	0	250	23	2519	10.00	100.76
P8_04_UKDC01-9080-MHE	2020-05-11 11:05	80	0	80	8	Running	21	51	0	2268	22	719	108.00	34.24
P8_05_UKDC01-9080-MHE	2020-05-11 11:05	80	0	80	8	Running	30	42	0	30	21	898	1.00	29.93
P8_06_UKDC01-9080-MHE	2020-05-11 11:05	80	0	80	8	Running	23	49	0	1012	21	1908	44.00	82.96
P8_07_UKDC01-9080-MHE	2020-05-11 11:05	80	0	80	8	Running	34	38	0	340	23	1522	10.00	44.76
P8_08_UKDC01-9080-MHE	2020-05-11 11:05	80	0	80	8	Running	27	45	0	1188	21	1547	44.00	57.30
P8_05_UKDC02-9080-MHE	2020-05-11 11:05	80	0	80	8	Available	0	2	0	0	0	1260	0.00	0.00
Total	2020-05-11 11:05	-	-	-	-	-	210	368	0	-	-	-	-	-

Welcome to the IBM POWER Capacity on Demand Dashboard

Overview | Capacity ≡ | Processor ≡ | Memory

Capacity on Demand History Information

All Logs | Resource | Access Key

System	Date	Time	Code	Activity	Activity Details	Additional Information
P8_05_UKDC02-9080-MHE	04/22/2020	13:42:57	HSCL034B	Mobile CoD memory set during power on	GB of memory: 847.	
P8_05_UKDC02-9080-MHE	04/22/2020	13:42:56	HSCL034A	Mobile CoD processors set during power on	number of processors: 30.	
P8_05_UKDC02-9080-MHE	04/22/2020	13:42:54	HSCL034B	Mobile CoD memory set during power on	GB of memory: 847.	
P8_05_UKDC02-9080-MHE	04/22/2020	13:42:53	HSCL034A	Mobile CoD processors set during power on	number of processors: 30.	
P8_08_UKDC01-9080-MHE	03/26/2020	11:57:04	HSCL032B	3 additional processor days will be charged for the hours remaining in the current processor day for the additional processors.		
P8_08_UKDC01-9080-MHE	03/26/2020	11:57:03	HSCL0329	On/Off CoD processors changed	number of processors: 27	number of days: 90.
P8_06_UKDC01-9080-MHE	03/26/2020	11:55:08	HSCL032B	2 additional processor days will be charged for the hours remaining in the current processor day for the additional processors.		
P8_06_UKDC01-9080-MHE	03/26/2020	11:55:06	HSCL0329	On/Off CoD processors changed	number of processors: 23	number of days: 90.
P8_05_UKDC02-9080-MHE	03/06/2020	13:50:48	HSCL0346	Mobile CoD processors activated	number of processors: 10	total number of Mobile CoD processors on server: 30.
P8_04_UKDC01-9080-MHE	02/28/2020	20:03:53	HSCL032B	1 additional processor days will be charged for the hours remaining in the current processor day for the additional processors.		
P8_04_UKDC01-9080-MHE	02/28/2020	20:03:52	HSCL0329	On/Off CoD processors changed	number of processors: 21	number of days: 180.
P8_02_UKDC01-9080-MHE	02/28/2020	20:03:14	HSCL032B	4 additional processor days will be charged for the hours remaining in the current processor day for the additional processors.		
P8_02_UKDC01-9080-MHE	02/28/2020	20:03:13	HSCL0329	On/Off CoD processors changed	number of processors: 24	number of days: 180.
P8_07_UKDC01-9080-MHE	02/21/2020	08:28:07	HSCL0329	On/Off CoD processors changed	number of processors: 34	number of days: 90.
P8_03_UKDC01-9080-MHE	02/21/2020	08:24:19	HSCL0329	On/Off CoD processors changed	number of processors: 25	number of days: 90.
P8_01_UKDC01-9080-MHE	02/21/2020	08:20:30	HSCL032B	1 additional processor days will be charged for the hours remaining in the current processor day for the additional processors.		
P8_01_UKDC01-9080-MHE	02/21/2020	08:20:27	HSCL0329	On/Off CoD processors changed	number of processors: 26	number of days: 90.

CoD Dashboard - Processor Pool Views

Resource view for processor pool utilisation

- The CPU activations are displayed alongside the current, day & day peak usage, helping to identify issues of over assignment of processors compared to the actual utilisation
- Simplified breakout charts of specific data

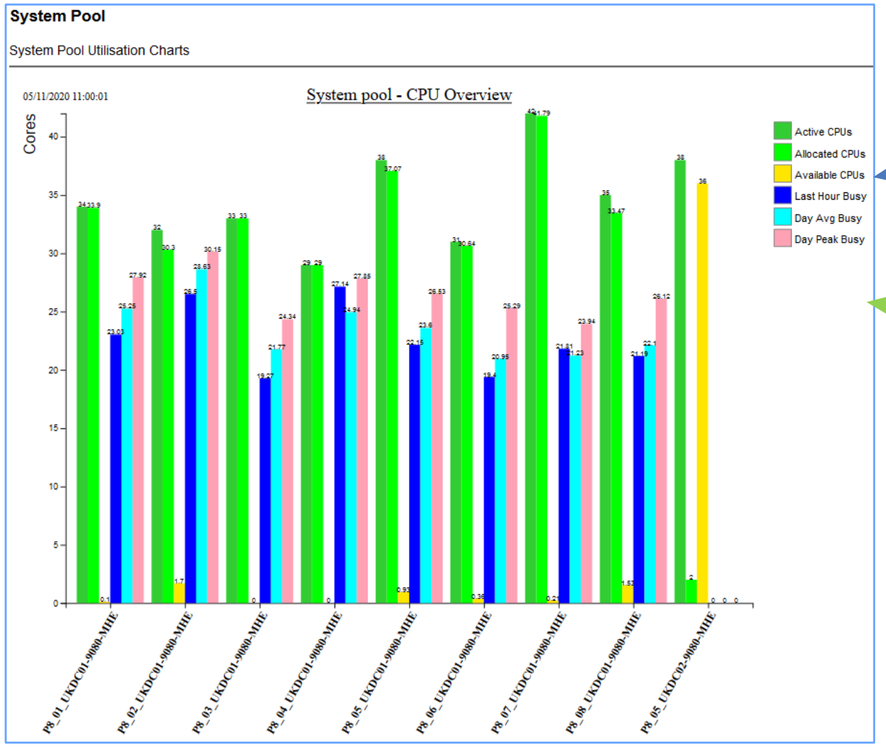

Utilisation

Breakout Charts

Configuration

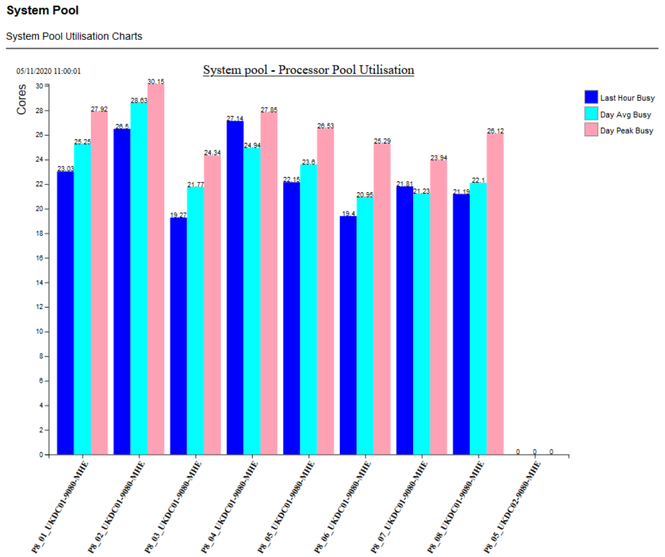

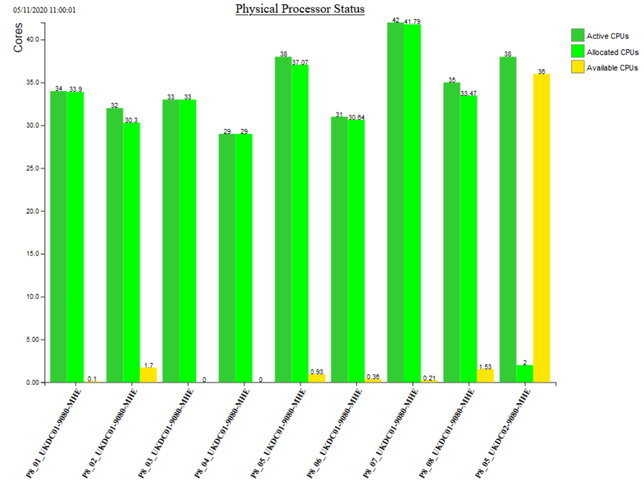

CoD Dashboard - Processor Pool Views...

- Historical charts help determine if current values are typical and show utilisation trends

  - 30 day stacked view allows tracking of entire estate

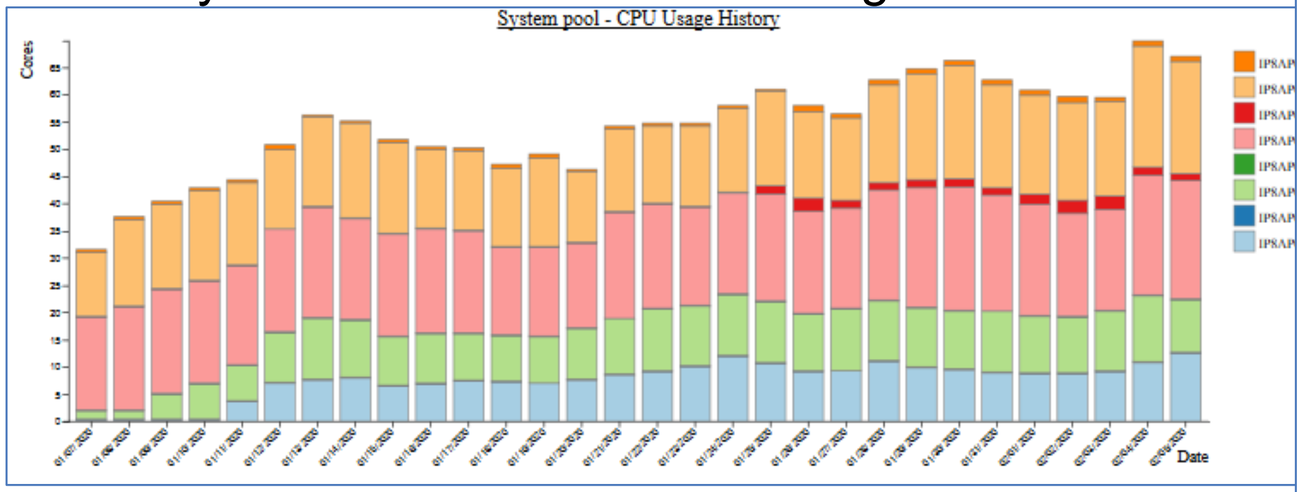

  - Seven Day history for quick trend checking

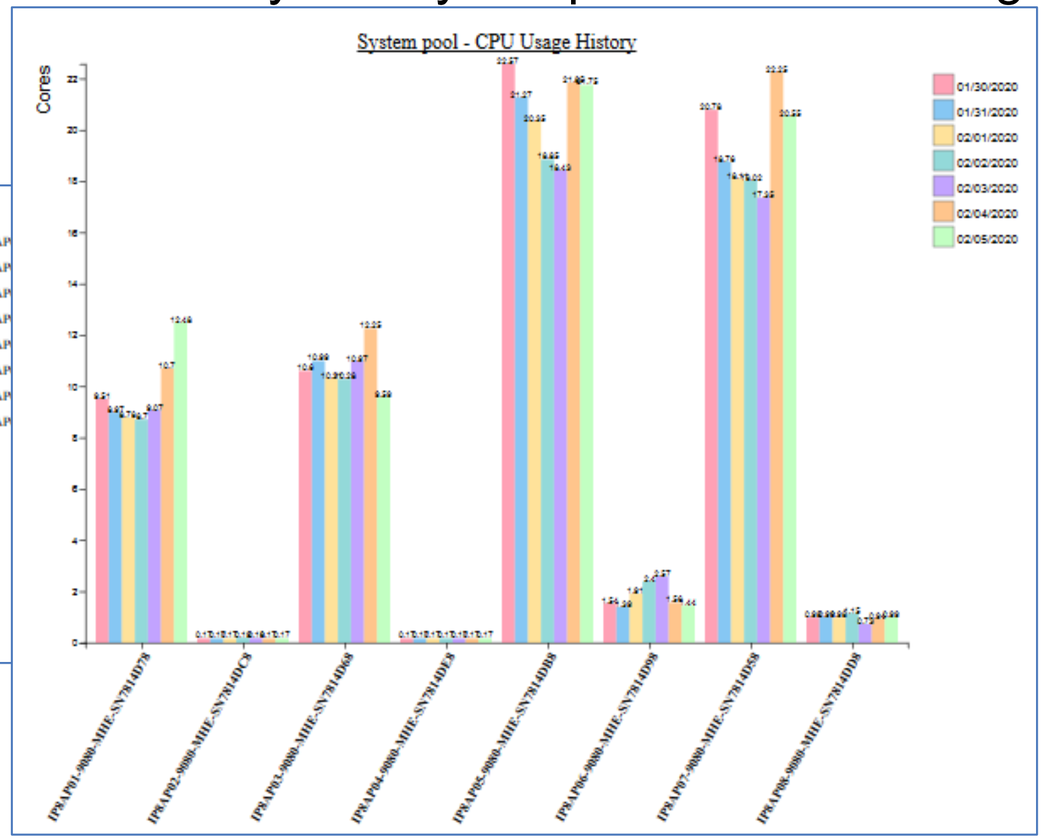

CoDD Tool Enablement

CoDD Offering

- Provided through a Systems Lab Services Engagement using a Power to Cloud voucher

CoDD Operating System requirements

- AIX or Linux environment with SSH access to HMC(s)
  - Non-privileged user accounts used to gather data from the HMCs
  - Tool uses standard Unix components of the Host OS
  - Minimal storage required for the scripts, dashboard files and data
- HTTP server (Apache/ngnix) required to display dashboard [1.]

CoDD Data

- HMC data is collected and processed by the tool to produce the dashboard
- Collected data and processed output are available in CSV files for use by customer tooling
- CoD history is collated each time the CoDD scripts run, all other data is refreshed hourly
- The number of pre-purchased resource days can be added to the tool to show their burn rate [2.]

1. Dashboard display requires an browser with internet access, or local copy of “d3” library
2. The number of pre-purchased resource days can be checked on the Entitled Systems Support (ESS) site, and then must be manually added to the CoDD Tool.

PowerView

Point in time Health Check Dashboard

PowerView

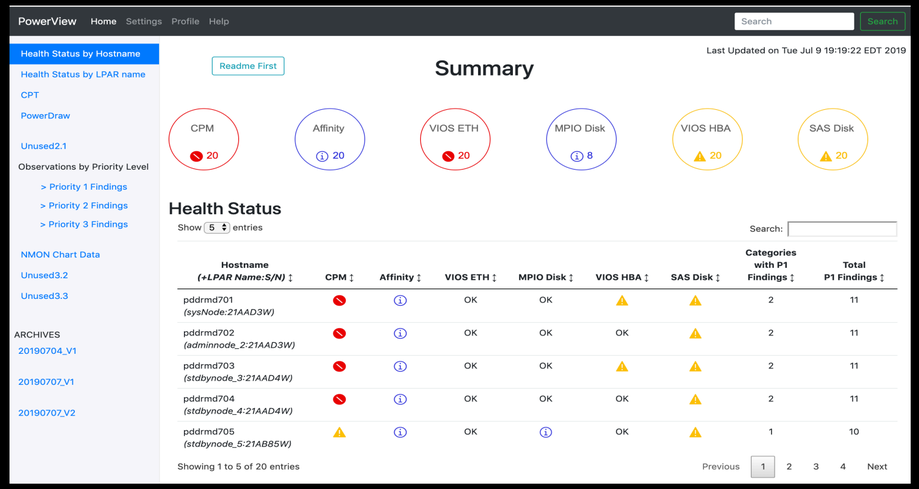

Daily Data Collection (AIX/VIOS)

- ✓ Point in time ‘agentless’ data collection scheduled via crontab.
- ✓ Collects configuration and performance data via standard AIX commands.
- ✓ ALL commands are light weight and non intrusive

Analyze (Linux/AIX)

- ✓ Repository on AIX or Linux
- ✓ Apache server for dashboard creation; optional NFS share for collection phase
- ✓ CSV file, excel file and html file creation are performed using AIX/VIOS collections

Present (HTML/excel)

- ✓ View assessment results in familiar dashboard format
- ✓ Navigation panel (left) presents findings by topic, both table and graphical formats
- ✓ Various of predefined health checks can be enhanced by adding custom checks

PowerView Overview

What is PowerView

- It‘s a point in time based health check

Main Components

- Data Collection (AIX/VIOS)
- Analyze (Linux/AIX)
- Present (HTML/excel)

Logical View

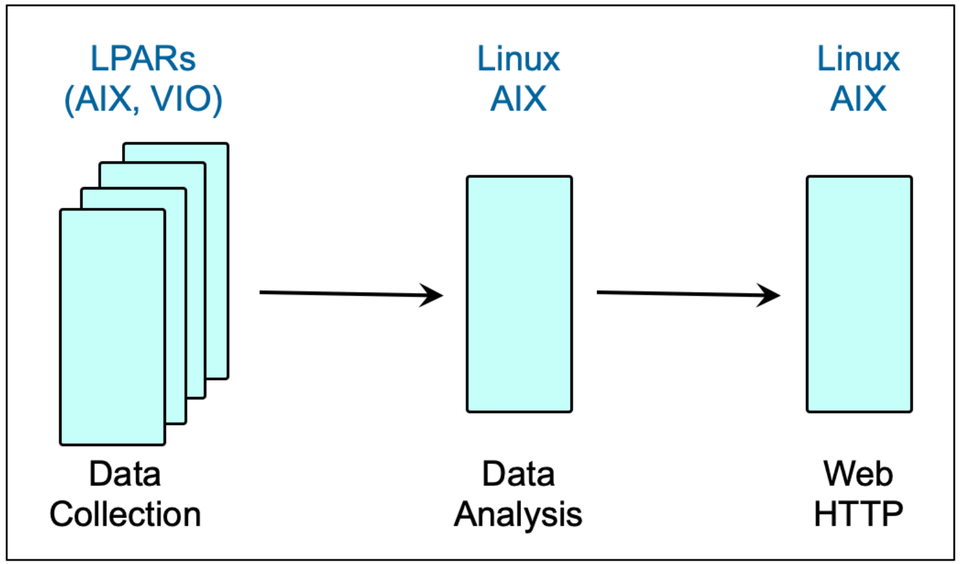

Setup Example

Using one system for both analysis and HTTP server

– Also hosts the executables and collected data

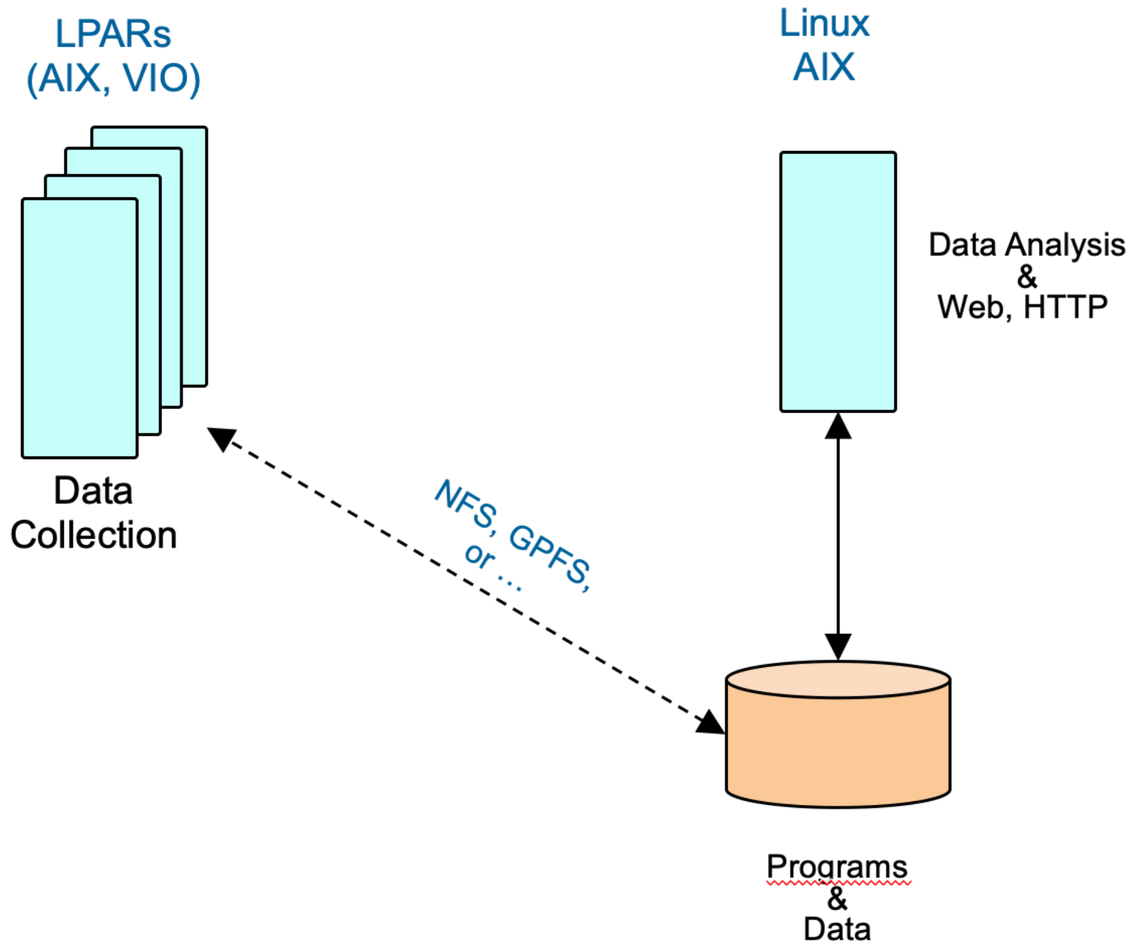

Setup Example

Using dedicated NFS server and one system for both analysis and HTTP server

- Also hosts the executables and collected data

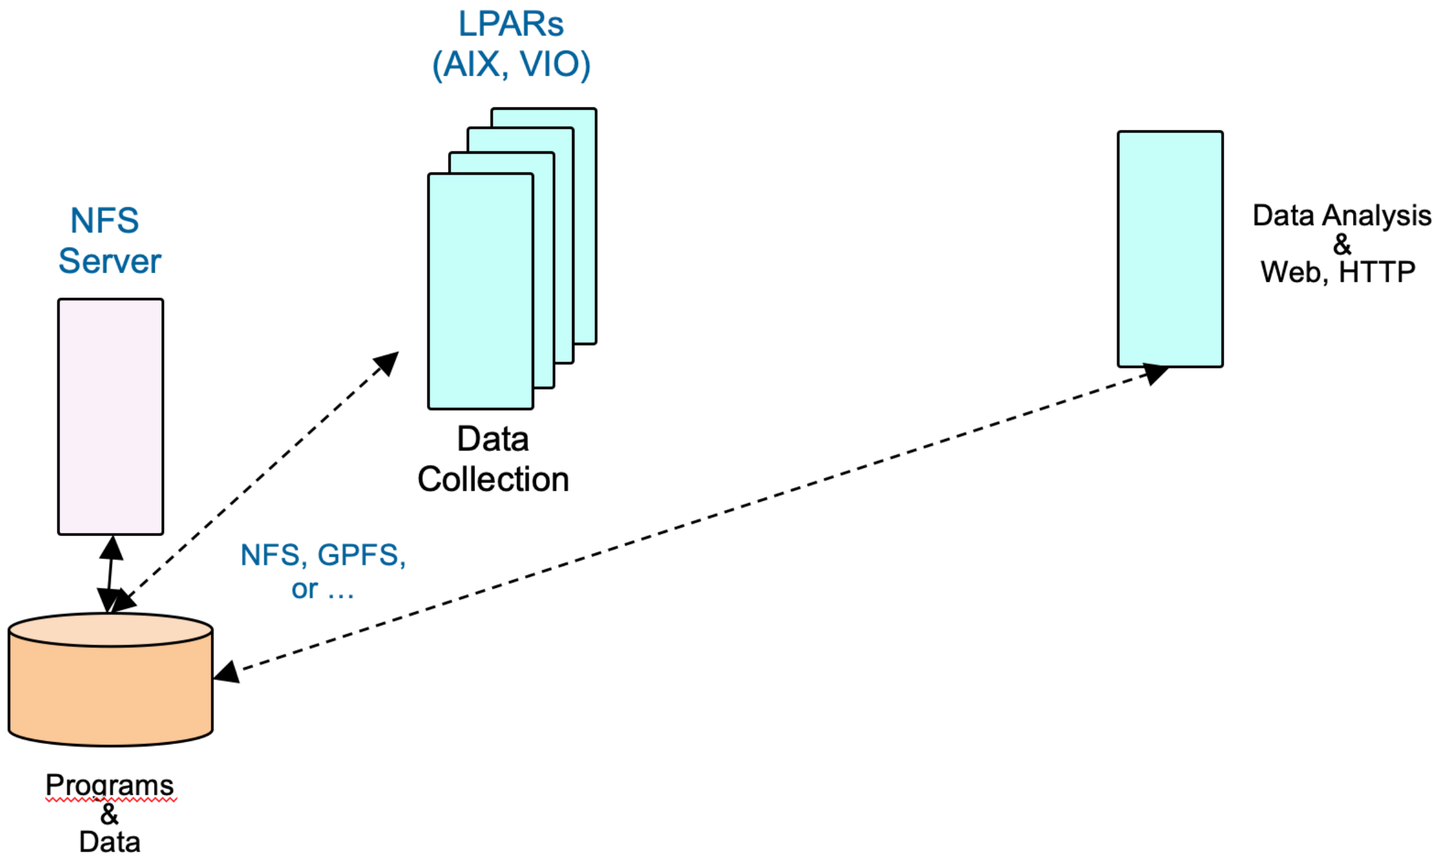

PowerView Dashboard

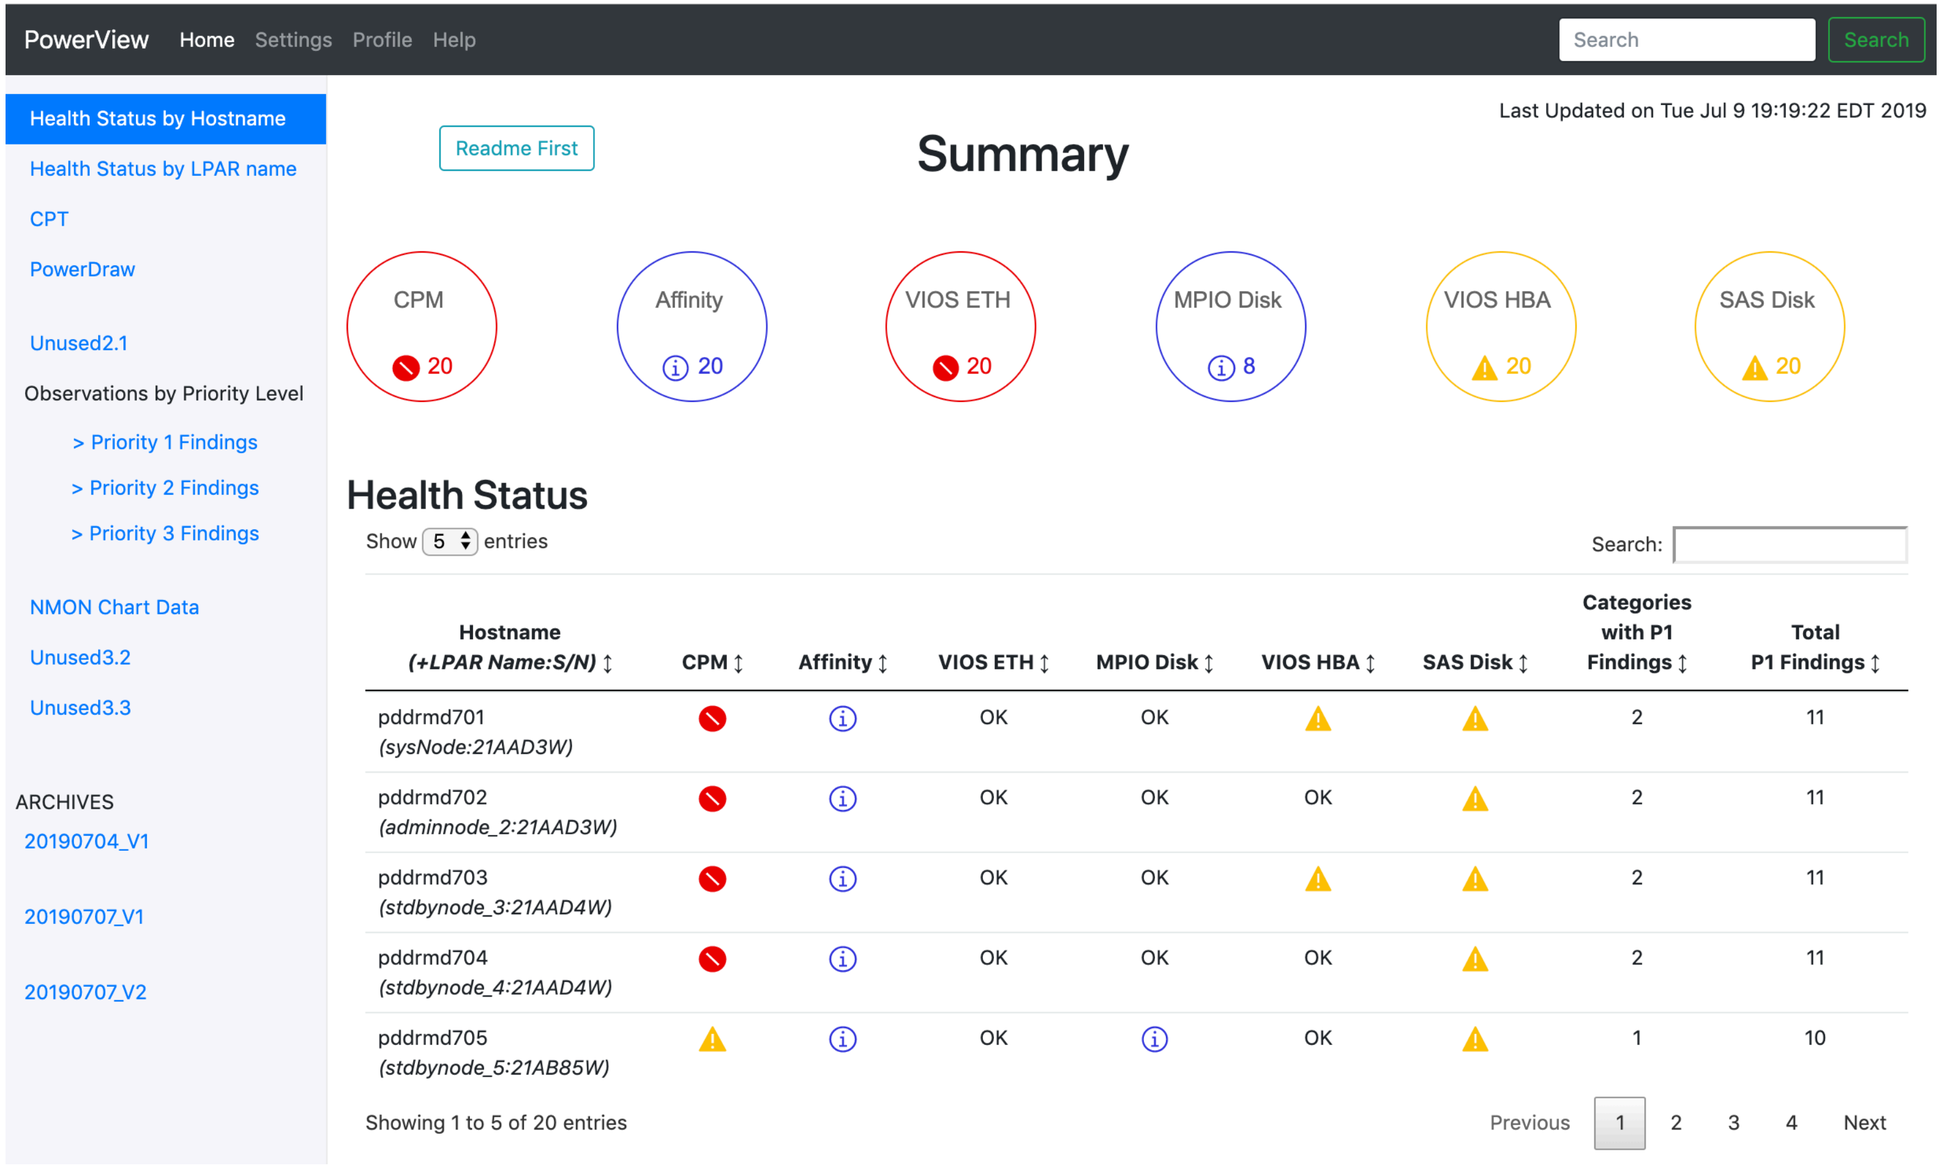

PowerDraw

Graphical view of Power Systems

PowerDraw

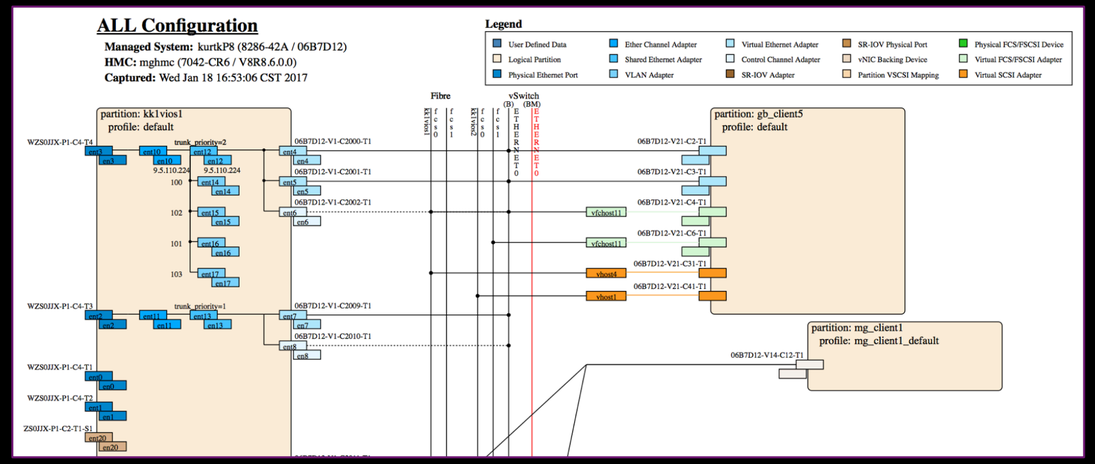

Visualize system virtualization design

Document components, linkages, parameters

Explore, analyze and debug

PowerDraw

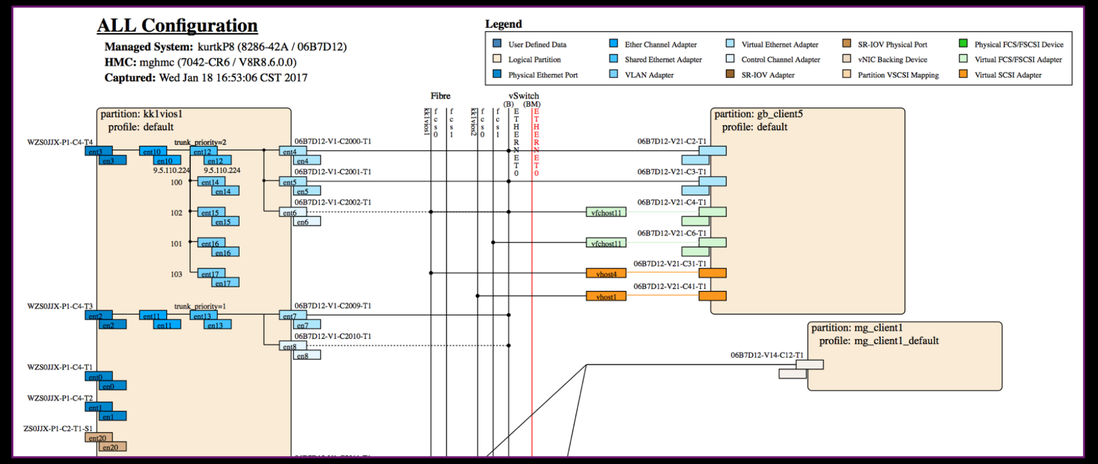

Business Challenge

- Unable to visualize configuration of a given managed system.
- Lacking ability to document key components in PowerVM design.
- Unable to graphically drill down to key elements within the PowerVM design of a managed system to query attributes or debug current problems or issues.

Solution

- PowerDraw which provides the ability to gather and display key configuration information within the managed system.
- Layered approach allows for viewing details at all or different levels within the PowerVM design

Key Benefits

- Ability to view any level of configuration for documentation purposes.
- Ability to perform problem determination using a graphical approach with current configuration settings

PowerDraw – at a glance

Drawframe

- **Capture and draw from existing design**
  - capture all/any managed systems and/or partitions from managing HMC
  - draw all/any managed systems and/or partitions from an HMC
  - merge user defined data so PowerDraw becomes launch point for any system related data

Views Generated

- **Network**
- **Storage**
- **SR-IOV / vNIC**
- **All**

Implementation

- **SVG (Scalable Vector Graphics), JavaScript**

Browser

- **Consumer of draw files produced by PowerDraw**
- **Displays detailed data utilizing onclick, onmouseover, onmouseout**

PowerDraw

- **Visualize** system virtualization design
- **Document** components, linkages, parameters
- **Explore,** analyze and debug

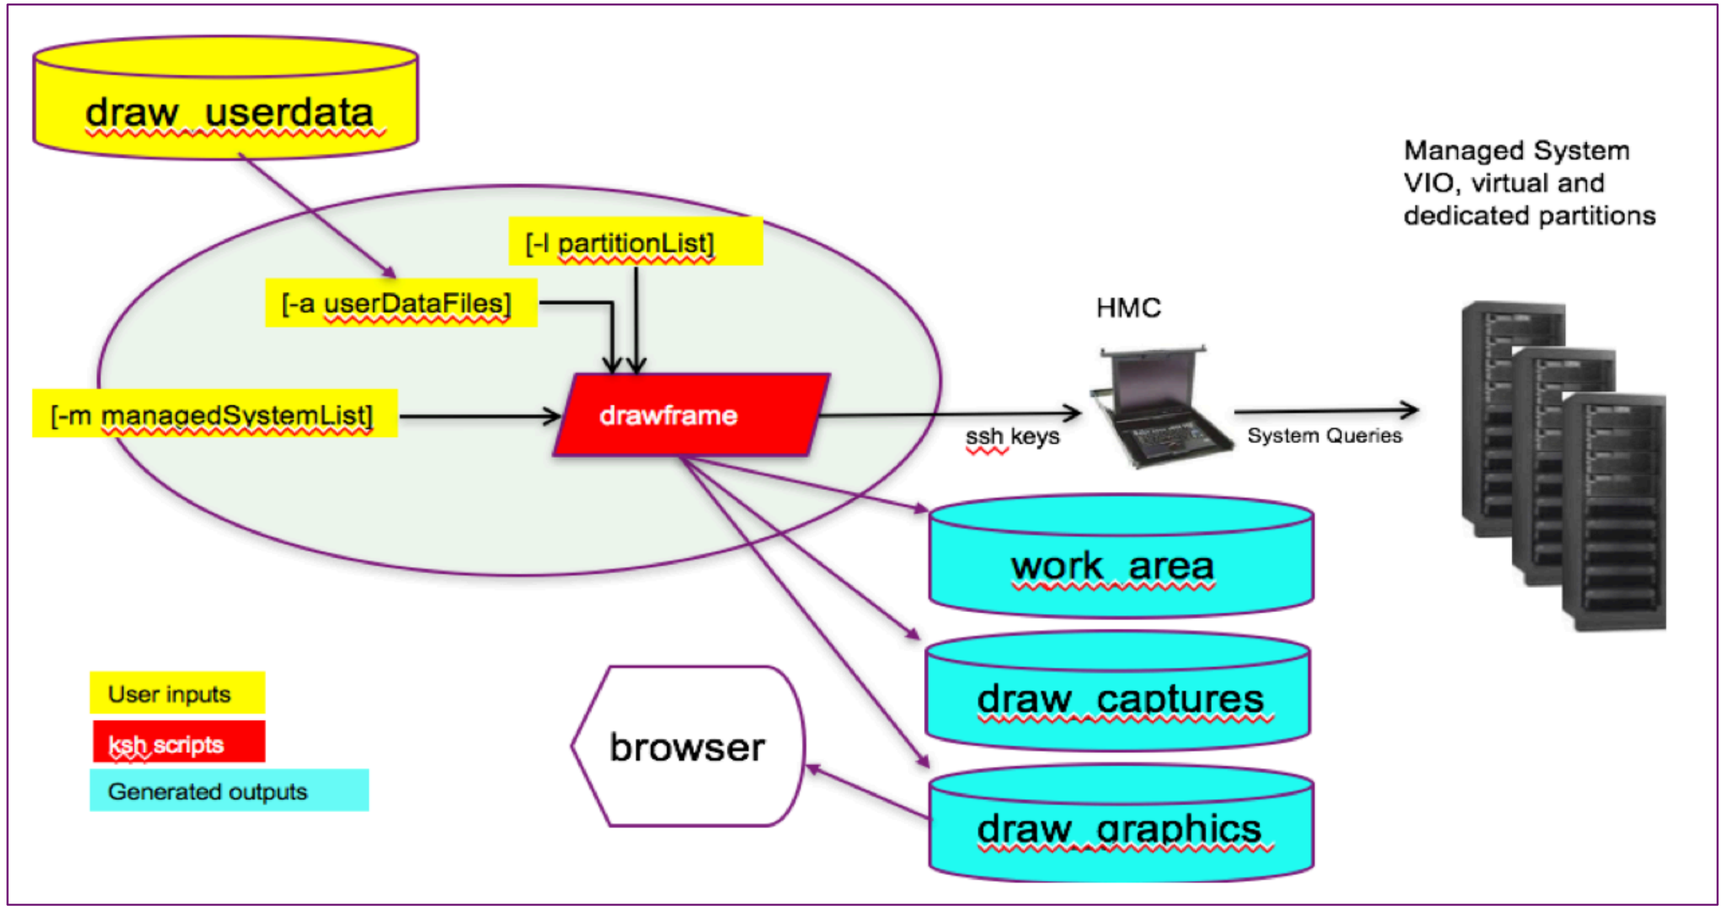

PowerDraw – NETWORK view

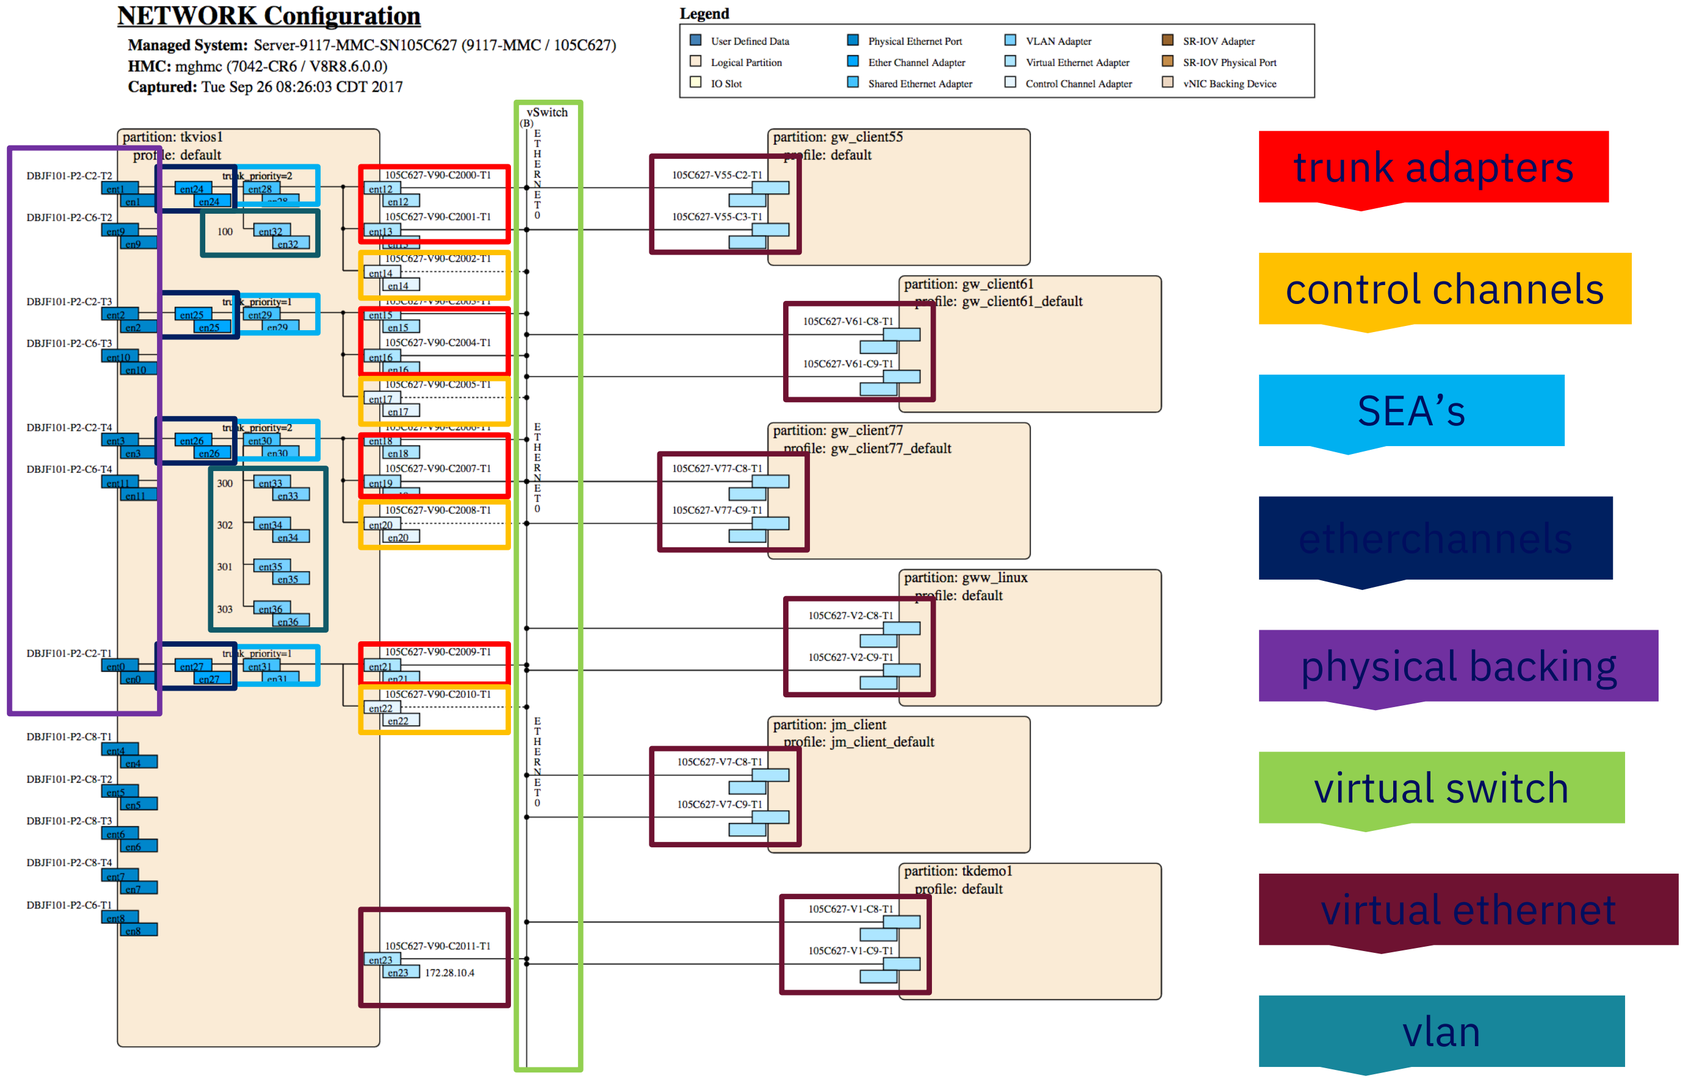

trunk adapters

control channels

SEA's

etherchannels

physical backing

virtual switch

virtual ethernet

vlan

PowerDraw – STORAGE view

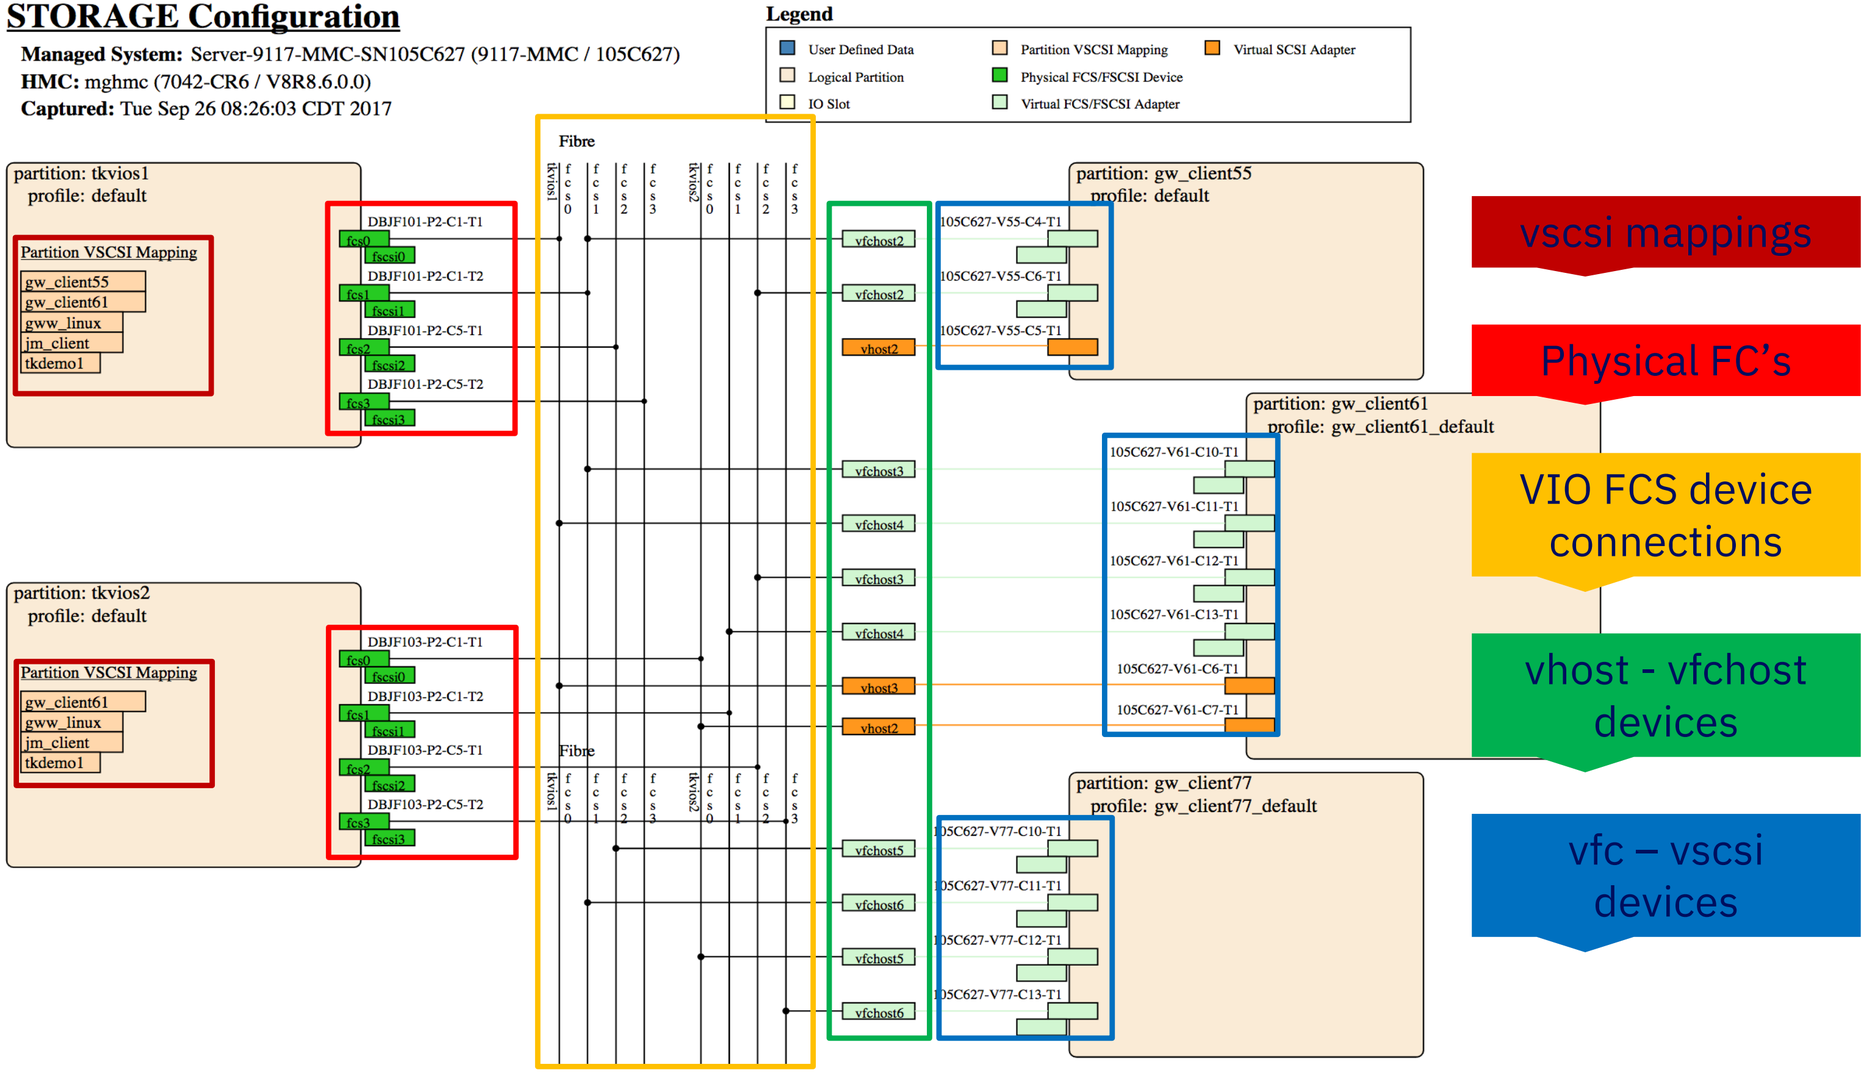

PowerDraw – VNIC view

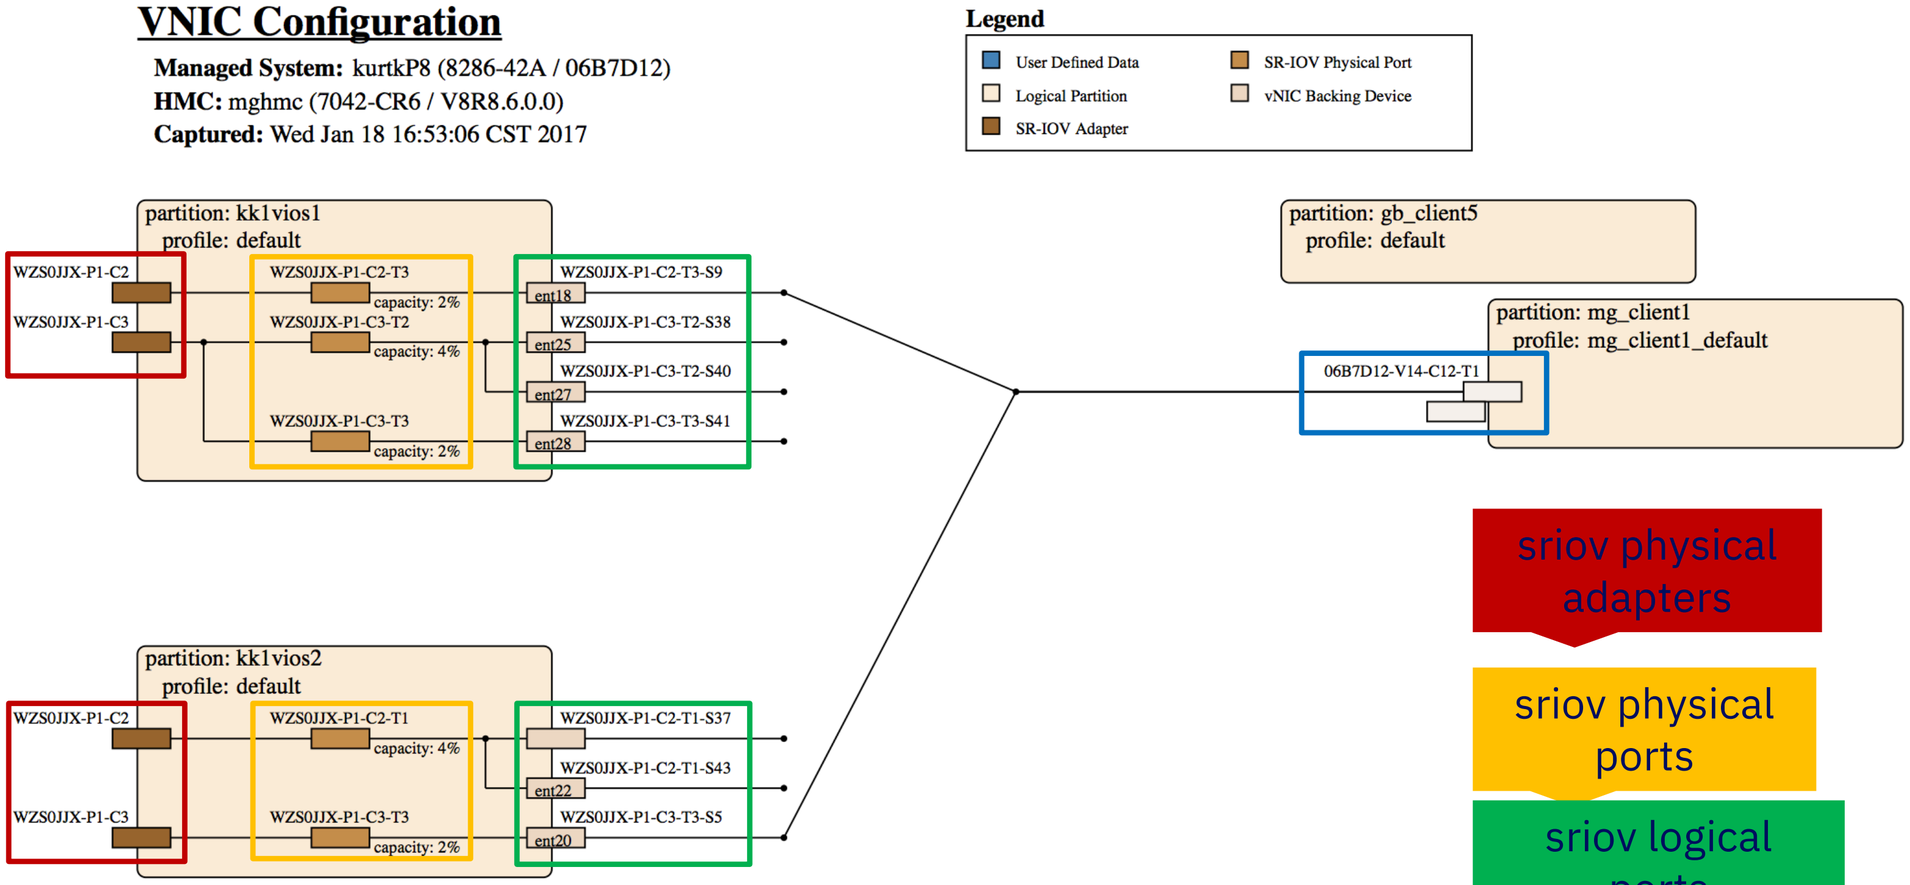

sriov physical adapters

sriov physical ports

sriov logical ports

vNIC backing device

PowerDraw

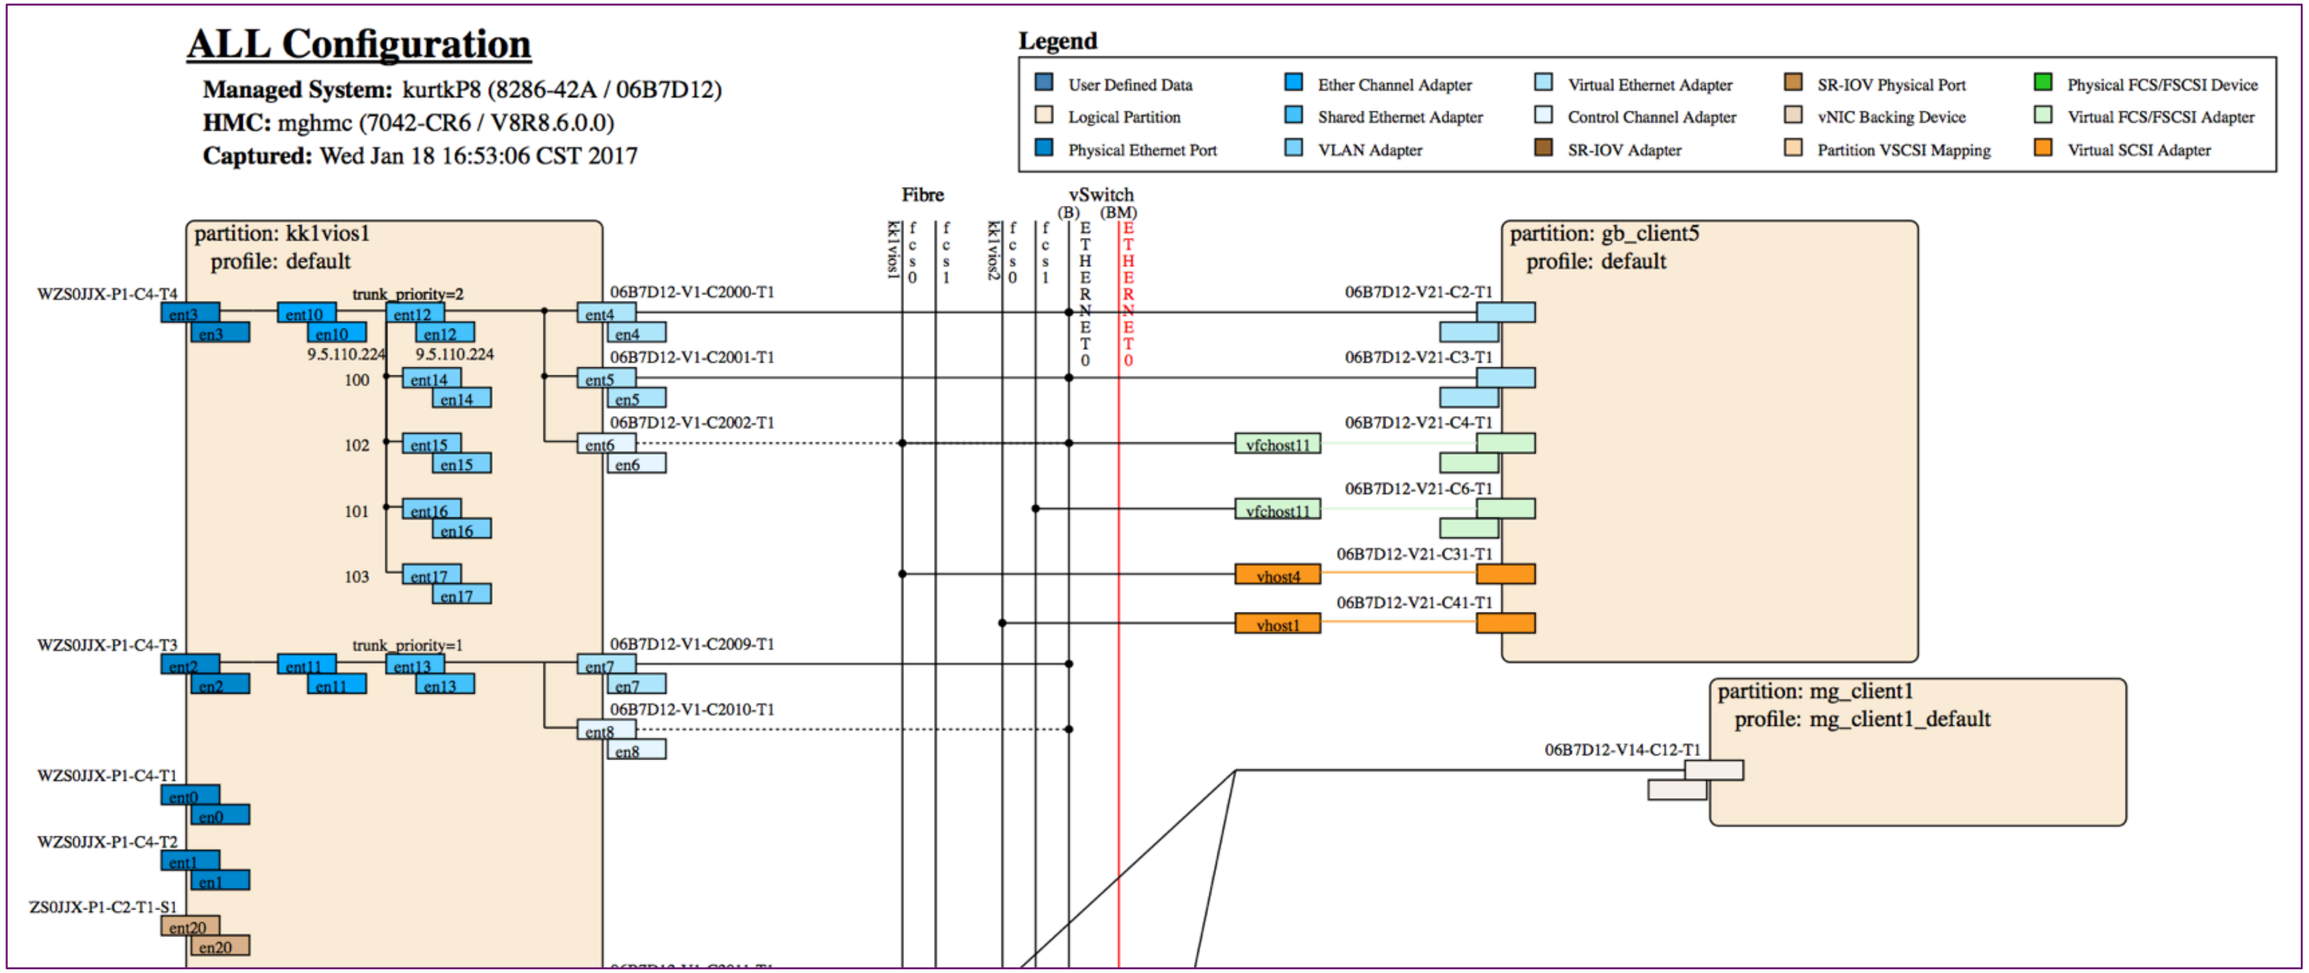

IBM Systems Lab Services Youtube channel - Demo:

PowerDraw - Graphical View of Power Systems https://www.youtube.com/watch?v=8P6q87bU5qo

ProMon

Strategic Monitoring

ProMon

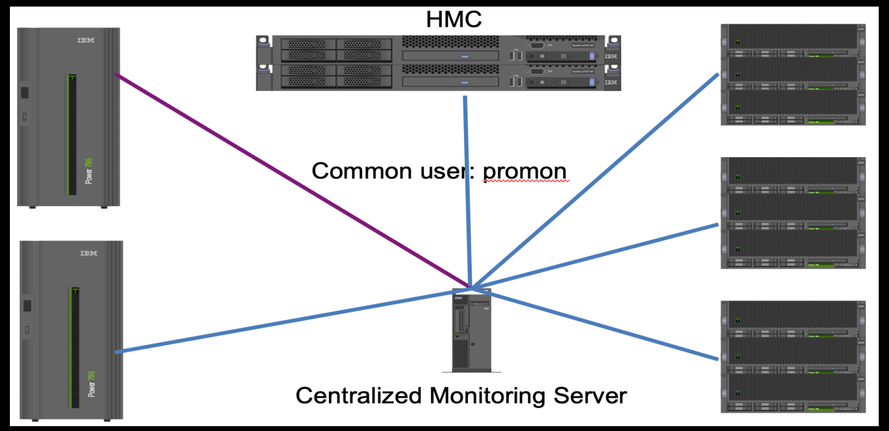

Why

- Your system is deployed perfectly, keep it that way with proactive monitoring.
- All systems tend toward chaos. Even with proper build and change procedures, people still make errors
- Redundancy built and properly verified during initial install, but it only protects against the first failure

Solution

- Introduce ProMon which runs daily audits to;
  - verify system configuration
  - alert system administrators against changes
  - and ensure the stability and health of the managed systems.

Key Benefits

- Prevent unplanned outages by leveraging a centralized auditing of multiple systems.
- Notify IT staff about looming problems before they become critical
- Provide real-time views of the :
  - » Resource status
  - » Resource allocation
  - » System configuration settings

ProMon

Strategic Monitoring

- Run daily audits across systems
- Custom and built-in predefined best practices and configuration checks
- Sends e-mail reports to system admins

Sample LPAR Report

LPAR Status Report Jul 21 2016
at 02:01:00 AM

Processing HMC unxhmcpa002
Collecting LPAR list from HMC unxhmcpa002
Processing HMC unxhmcpa003
Collecting LPAR list from HMC unxhmcpa003

Processing Frame florida
- Processing Frame mississippi
-
- lpar001-NoLPM, Unable to audit LPAR due to failure of ssh
- …
- lpar2, Check rootvg on lpar2 for closed/stale partitions
- lpar2, MPIO policy error. Fewer than 4 path to disk hdisk0
- lpar2, MPIO policy error. Fewer than 4 path to disk hdisk1
- lpar2, Shortage of DMA Resources found for fcs2
- lpar2, Consider changing command elements on fcs2 from 500 to 1024
- lpar2, Consider changing transfer size on fcs2 from 0x100000 to 0x200000
- lpar2, OS inconsistencies found on lpar2
- ….
- lpar3, Check rootvg on lpar3 for closed/stale partitions
- lpar3, Health Check interval policy error on hdisk12
- lpar3, Queue depth policy error for hdisk12
- lpar3, Health Check interval policy error on hdisk13
- lpar3, Queue depth policy error for hdisk13
- ….
- lpar4, TSM may not be running
- fcs and vscsi adapters both exist on lpar4
- …..
- lpar5, Check rootvg on lpar5 for closed/stale partitions
- lpar5, Reserve lock set to yes for hdiskpower0
- lpar5, Reserve lock set to yes for hdiskpower2
- lpar5, Reserve lock set to yes for hdiskpower3
- lpar5, Reserve lock set to yes for hdiskpower4
- lpar5, Shortage of DMA Resources found for fcs0

Sample VIOS report

Starting VIOS audits at 0257 on 052013

atl750dvio01
There is/are 2 configured Shared Ethernet Adapters on atl750dvio01 functioning normally
For ent10 there were 2 MB sent and 6 MB received with 36447 xmit errors and 0 receive errors
Maximum Queue Depth was 4 packets with no overflow errors
For ent11 there were 0 MB sent and 9 MB received with 0 xmit errors and 0 receive errors
Maximum Queue Depth was 2 packets with no overflow errors

Unmirrored LV found on atl750dvio01 rootvg hd5 boot N/A
Unmirrored LV found on atl750dvio01 rootvg hd6 paging N/A
Unmirrored LV found on atl750dvio01 rootvg paging00 paging N/A
Unmirrored LV found on atl750dvio01 rootvg hd8 jfs2log N/A

Actual VIOS report

sv65759 For ent14 there were 2 MB sent and 2 MB received with 0 xmit errors and 0 receive errors
sv65759 Maximum Transmit Queue was 0 packets with no overflow errors
sv65759, For ent18 there were 7 MB sent and 11 MB received with 55428 xmit errors and 0 receive errors
sv65759 Maximum Transmit Queue was 0 packets with no overflow errors

sv65759 Rootvg properly mirrored
sv65759, Dynamic tracking not set on fscsi0
sv65759, Fast_Fail attribute not set on fscsi0
sv65759, Dynamic tracking not set on fscsi1
sv65759, Fast_Fail attribute not set on fscsi1

Once a day audit against the entire environment

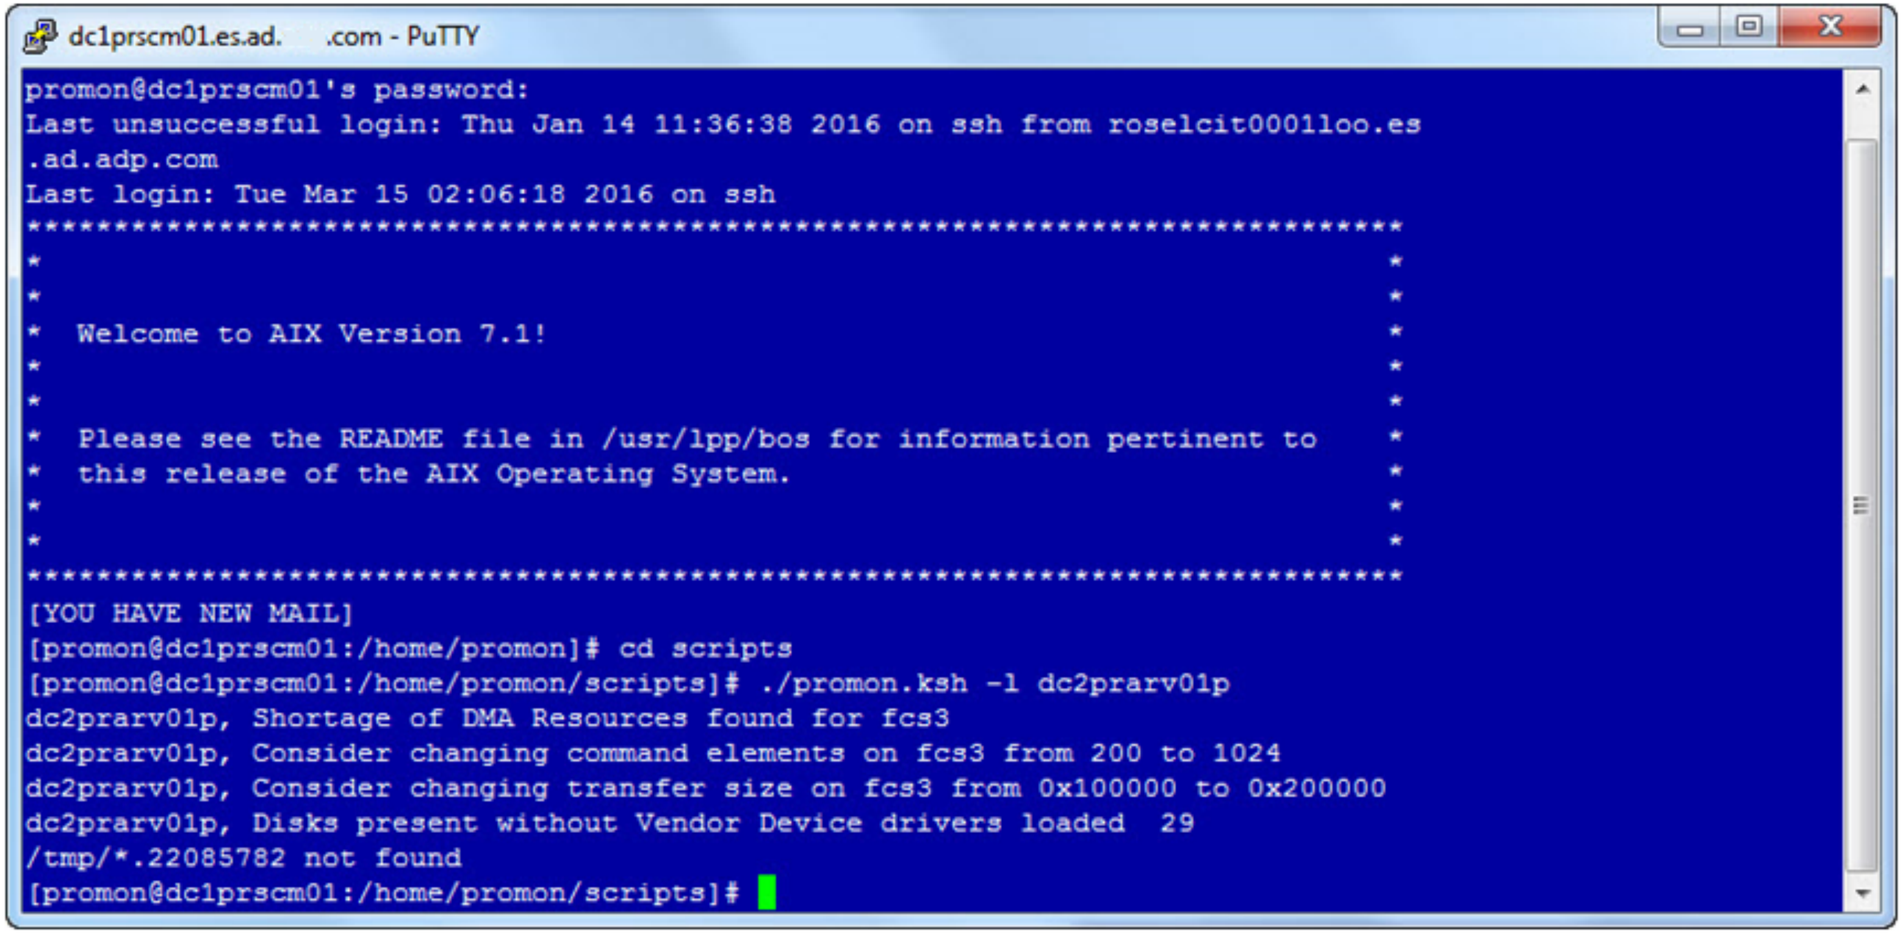

You can always run an audit against a single LPAR.

Capacity Planning Toolkit

Capacity Planning Toolkit

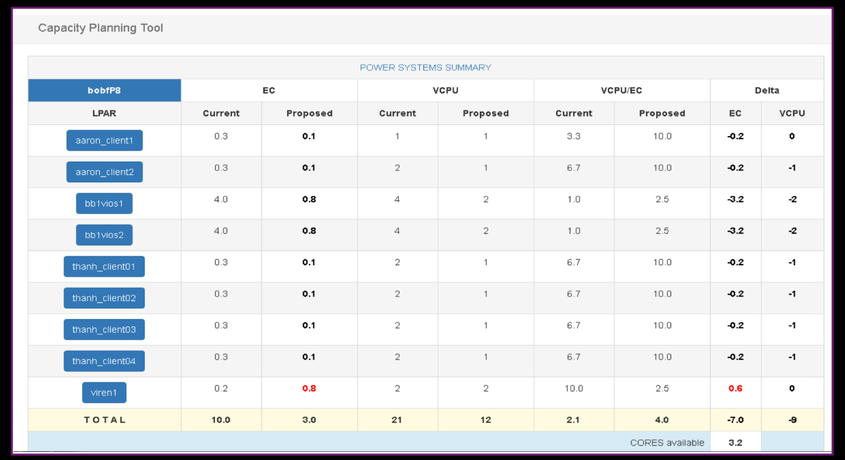

Capacity Planning Tool

POWER SYSTEMS SUMMARY

bobfP8	EC		VCPU		VCPU/EC		Delta	
LPAR	Current	Proposed	Current	Proposed	Current	Proposed	EC	VCPU
aaron_client1	0.3	**0.1**	1	1	3.3	10.0	**-0.2**	**0**
aaron_client2	0.3	**0.1**	2	1	6.7	10.0	**-0.2**	**-1**
bb1vios1	4.0	**0.8**	4	2	1.0	2.5	**-3.2**	**-2**
bb1vios2	4.0	**0.8**	4	2	1.0	2.5	**-3.2**	**-2**
thanh_client01	0.3	**0.1**	2	1	6.7	10.0	**-0.2**	**-1**
thanh_client02	0.3	**0.1**	2	1	6.7	10.0	**-0.2**	**-1**
thanh_client03	0.3	**0.1**	2	1	6.7	10.0	**-0.2**	**-1**
thanh_client04	0.3	**0.1**	2	1	6.7	10.0	**-0.2**	**-1**
viren1	0.2	**0.8**	2	2	10.0	2.5	**0.6**	**0**
TOTAL	**10.0**	**3.0**	**21**	**12**	**2.1**	**4.0**	**-7.0**	**-9**
						CORES available	3.2	

Designed to give helpful insight to your Power Systems capacity and guide right-sizing based on statistical analysis and IBM best practice recommendations from the experience of IBM Lab Services Power Systems performance consultants.

Helps clients better understand their current usage and plan for growth

Give accurate sizing for migration of legacy hardware to Power9 systems.

Capacity Planning Toolkit

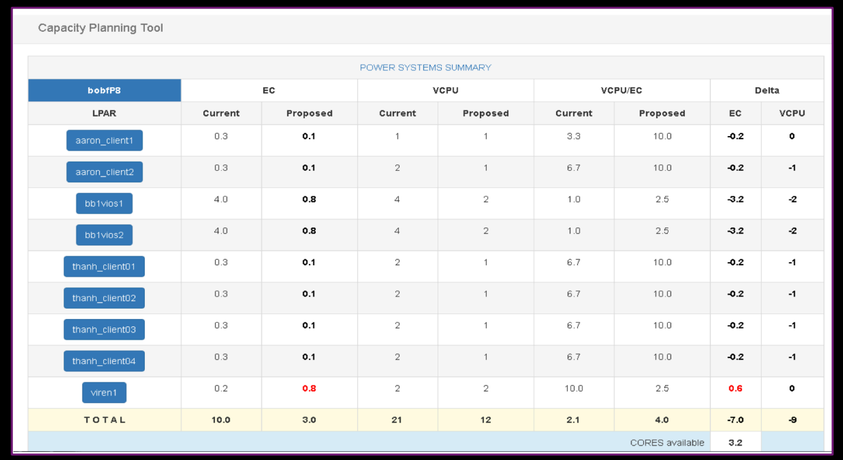

Capacity Planning Tool

POWER SYSTEMS SUMMARY

bobfP8	EC		VCPU		VCPU/EC		Delta	
LPAR	Current	Proposed	Current	Proposed	Current	Proposed	EC	VCPU
aaron_client1	0.3	**0.1**	1	1	3.3	10.0	**-0.2**	**0**
aaron_client2	0.3	**0.1**	2	1	6.7	10.0	**-0.2**	**-1**
bb1vios1	4.0	**0.8**	4	2	1.0	2.5	**-3.2**	**-2**
bb1vios2	4.0	**0.8**	4	2	1.0	2.5	**-3.2**	**-2**
thanh_client01	0.3	**0.1**	2	1	6.7	10.0	**-0.2**	**-1**
thanh_client02	0.3	**0.1**	2	1	6.7	10.0	**-0.2**	**-1**
thanh_client03	0.3	**0.1**	2	1	6.7	10.0	**-0.2**	**-1**
thanh_client04	0.3	**0.1**	2	1	6.7	10.0	**-0.2**	**-1**
viren1	0.2	**0.8**	2	2	10.0	2.5	**0.6**	**0**
TOTAL	**10.0**	**3.0**	**21**	**12**	**2.1**	**4.0**	**-7.0**	**-9**
						CORES available	3.2	

Analyze data from a variety of sources including NMON, LPAR2RRD / HMC **LPARUTIL**, ITM, BMC Patrol/Perceiver, HP Openview or Custom

Right-sizing recommendations backed by core Power performance team

Target specific frame utilizations

Incorporate specific business rules and SLAs with multiple rules based on service tiering

HA awareness, including failover capacity in sizing and Active/Active or Active/Standby

Averaging – In Graphic Detail

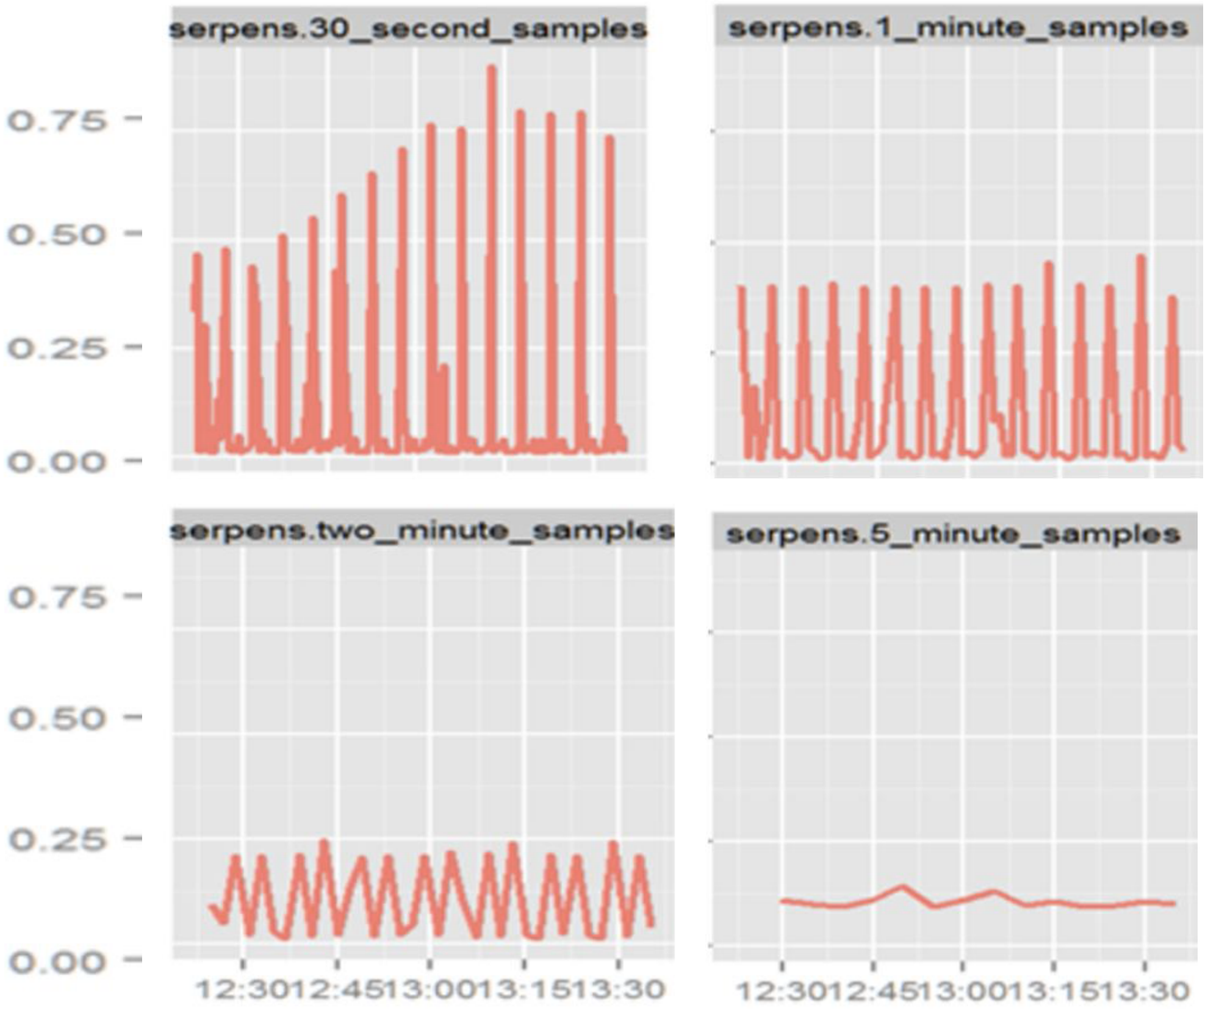

LPAR with
EC=1.0

CPT – sampling and duration:

- **Prefer frequent sampling**
  - Minutes or seconds – not hours or days
- **Want a duration that spans significant times in your business**
- **Calculate consumption by percentile**
  - Look at the frequency physical consumption exceeds certain values

CPT Strategy - Percentiles

Size partitions to some percentile of physc – e.g. 80th Percentile

- i.e. partition uses shared pool resources 20% of the time

"A percentile is a measure used in statistics indicating the value below which a given percentage of observations in a group of observations fall. "

Are you (and your customers) comfortable with sharing 20% of the time?

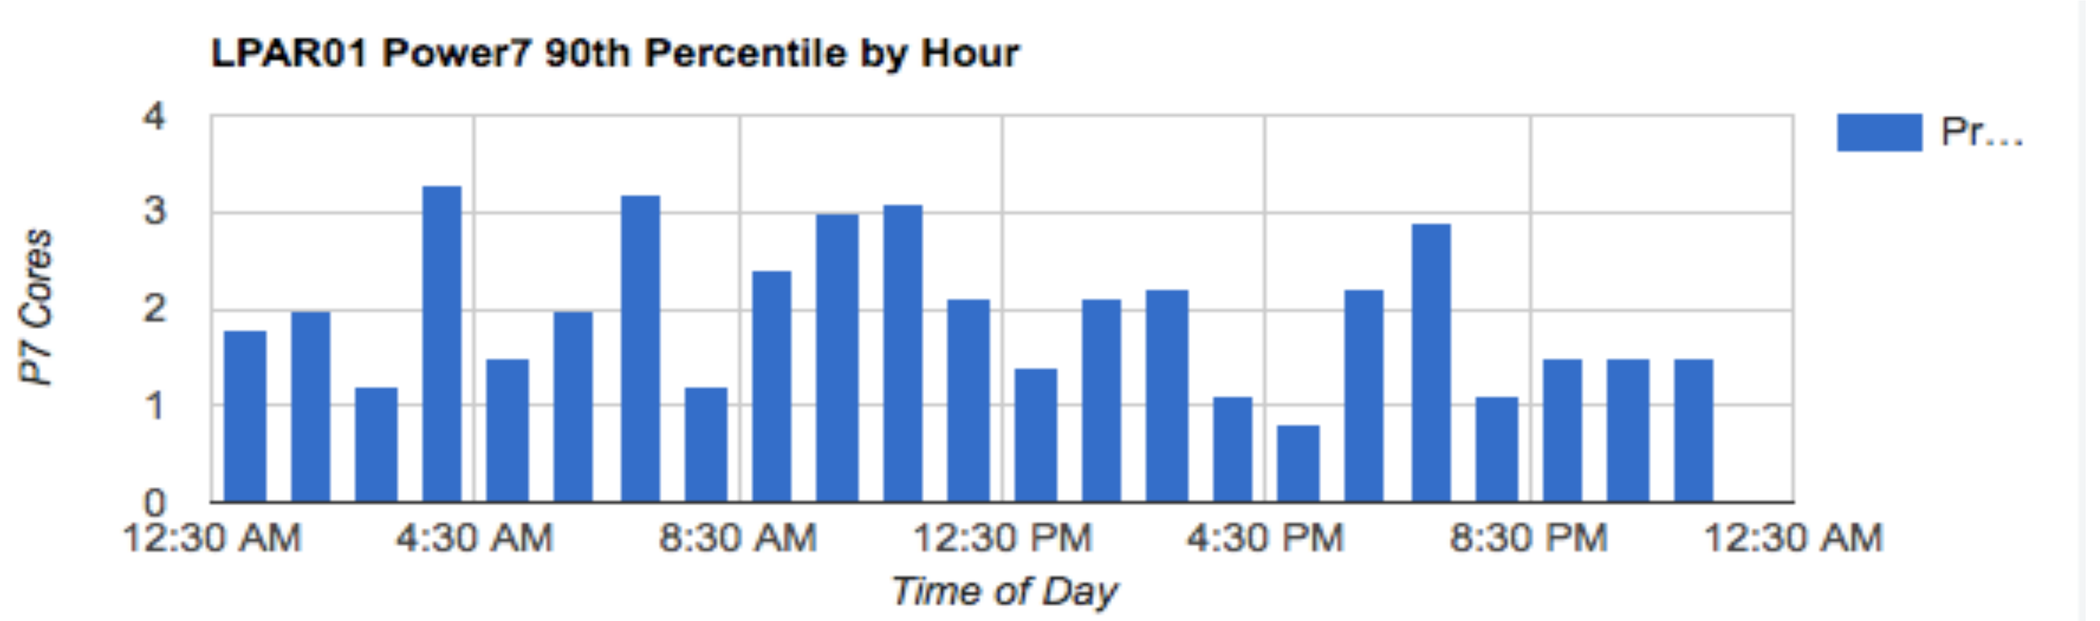
LPAR01 Power7 90th Percentile by Hour
P7 Cores
0
1
2
3
4
12:30 AM
4:30 AM
8:30 AM
12:30 PM
4:30 PM
8:30 PM
12:30 AM
Time of Day
Pr...

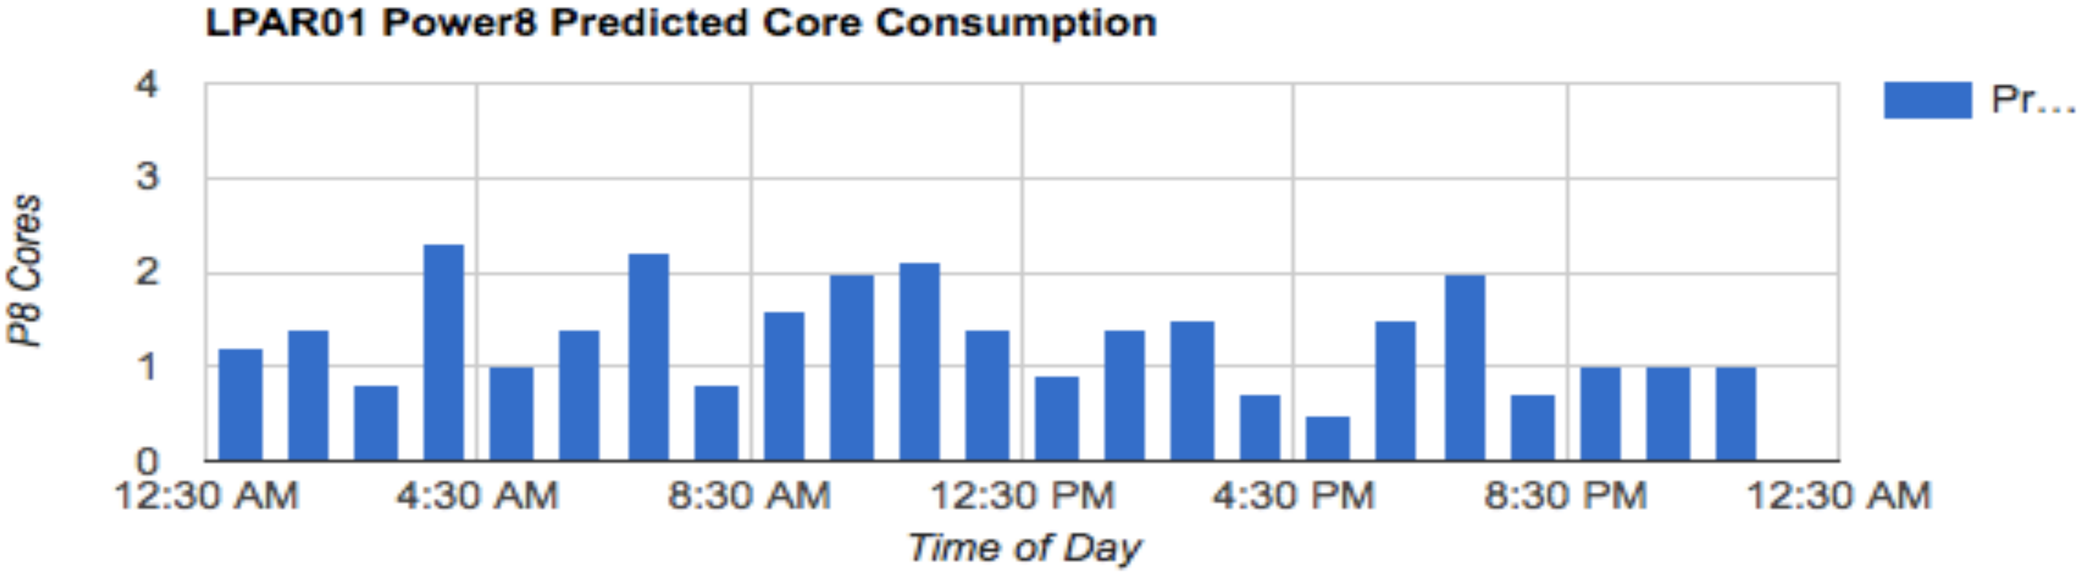
LPAR01 Power8 Predicted Core Consumption
P8 Cores
0
1
2
3
4
12:30 AM
4:30 AM
8:30 AM
12:30 PM
4:30 PM
8:30 PM
12:30 AM
Time of Day
Pr...

Physical CPU Consumption (a look under the hood)

```
hscroot@mghmc:~> lslparutil -m Server-9117-MMC-SN105C627 -s s
-r lpar --filter lpar_names=tk_client4 -n 5
-F time,lpar_name,time_cycles,capped_cycles,uncapped_cycles

08/19/2014 00:00:01,tk_client4,5712884416586958,3,210158055669066,6513452103872420
08/19/2014 00:01:01,tk_client4,5708460689894641,1,209983034697607,3,650780433652451
08/19/2014 00:02:01,tk_client4,5704036958144020,6,209804139410327,2,650193105704015
08/19/2014 00:03:01,tk_client4,5699613220126260,7,209624713801842,7,649610646232309
08/19/2014 00:04:01,tk_client4,5695189557264483,5,209450233526452,8,649050485284933
```

PHYSC =

(209624713801842**7**-209450233526452**8**) + (649610646232309**01**-649050485284933**13**)

(5699613220126260**7** - 5695189557264483**5**)

CPT – Capacity Planning Toolkit

Designed to be flexible, and handle a number of different input formats:

- NMON
- ITM
- BMC
- LPAR2RRD
- **HMC LPARUTIL**
- HMC PCM
- AIX VMSTAT
- CUSTOM

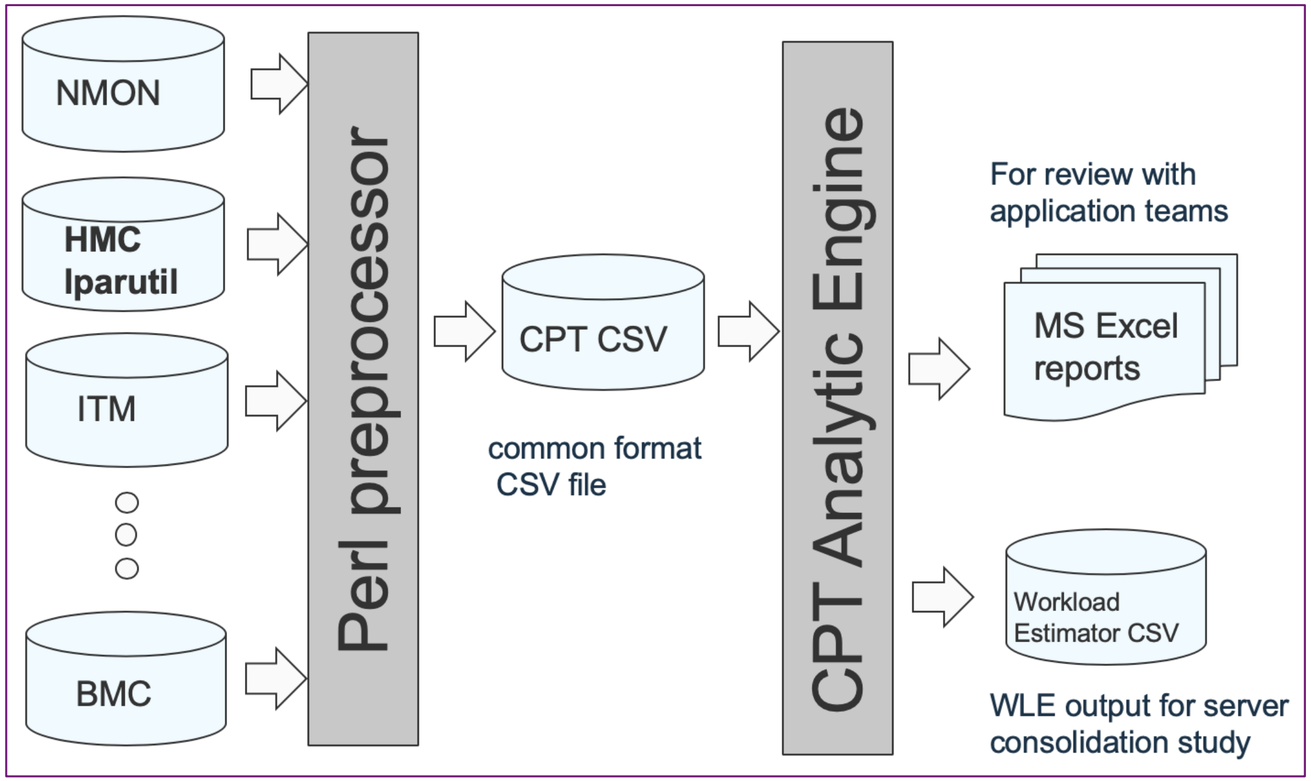

CPT – Capacity Planning Toolkit

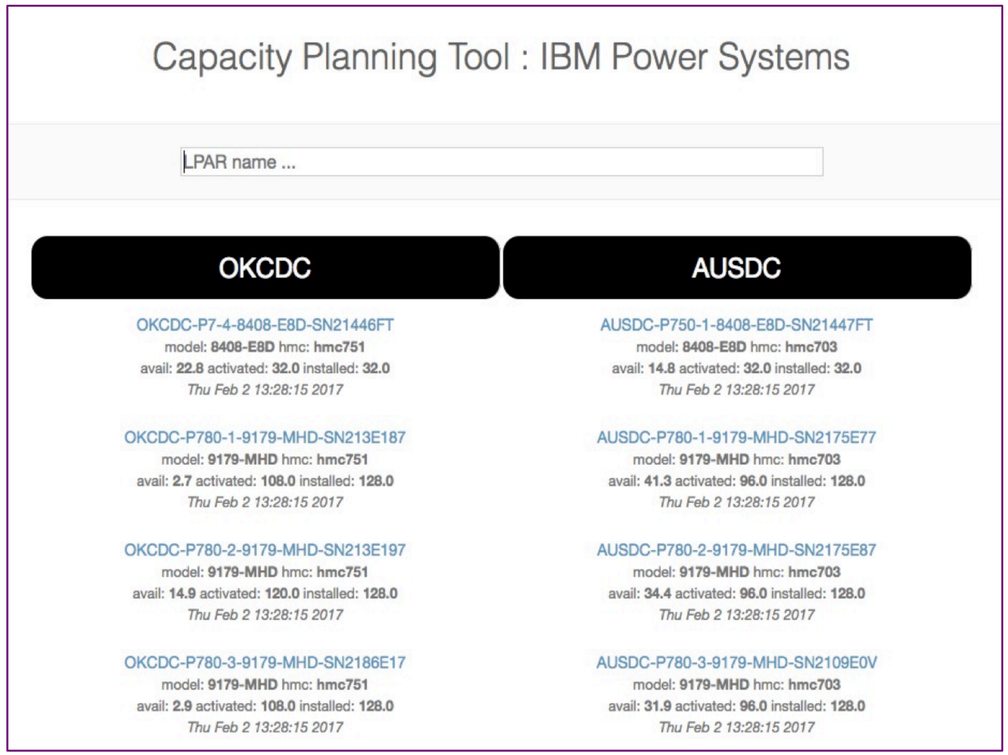

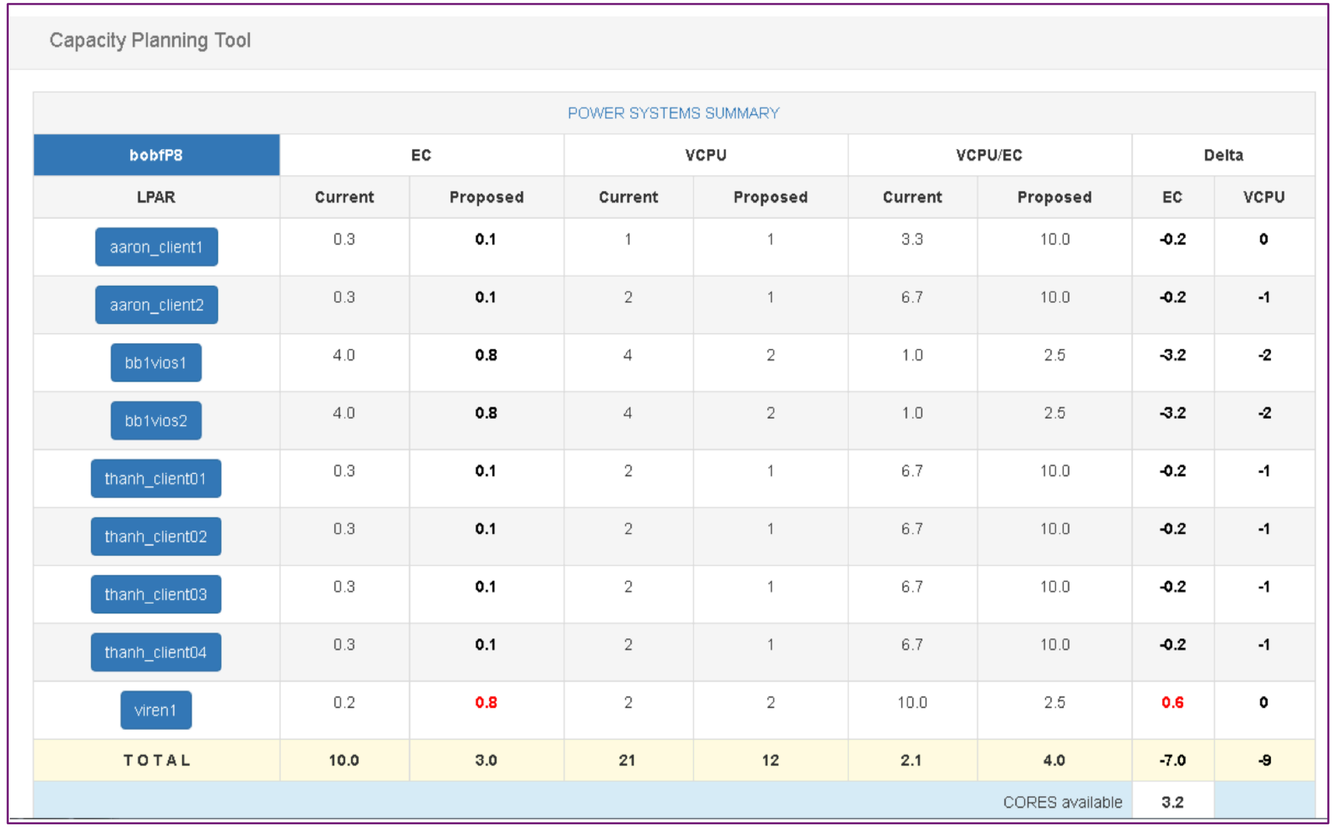

CPT – SCONN Server Consolidation

CPT - Server Consolidation

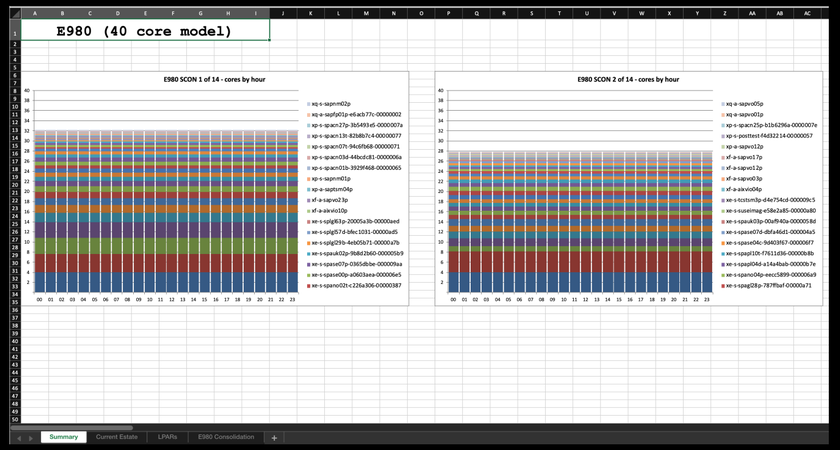

Give accurate sizing for migration of legacy hardware to Power9

Consolidate legacy Power systems using Rperf scaling

Typically used in a Migration Planning and Server Consolidation studies

Capacity Planning Tool – SCONN

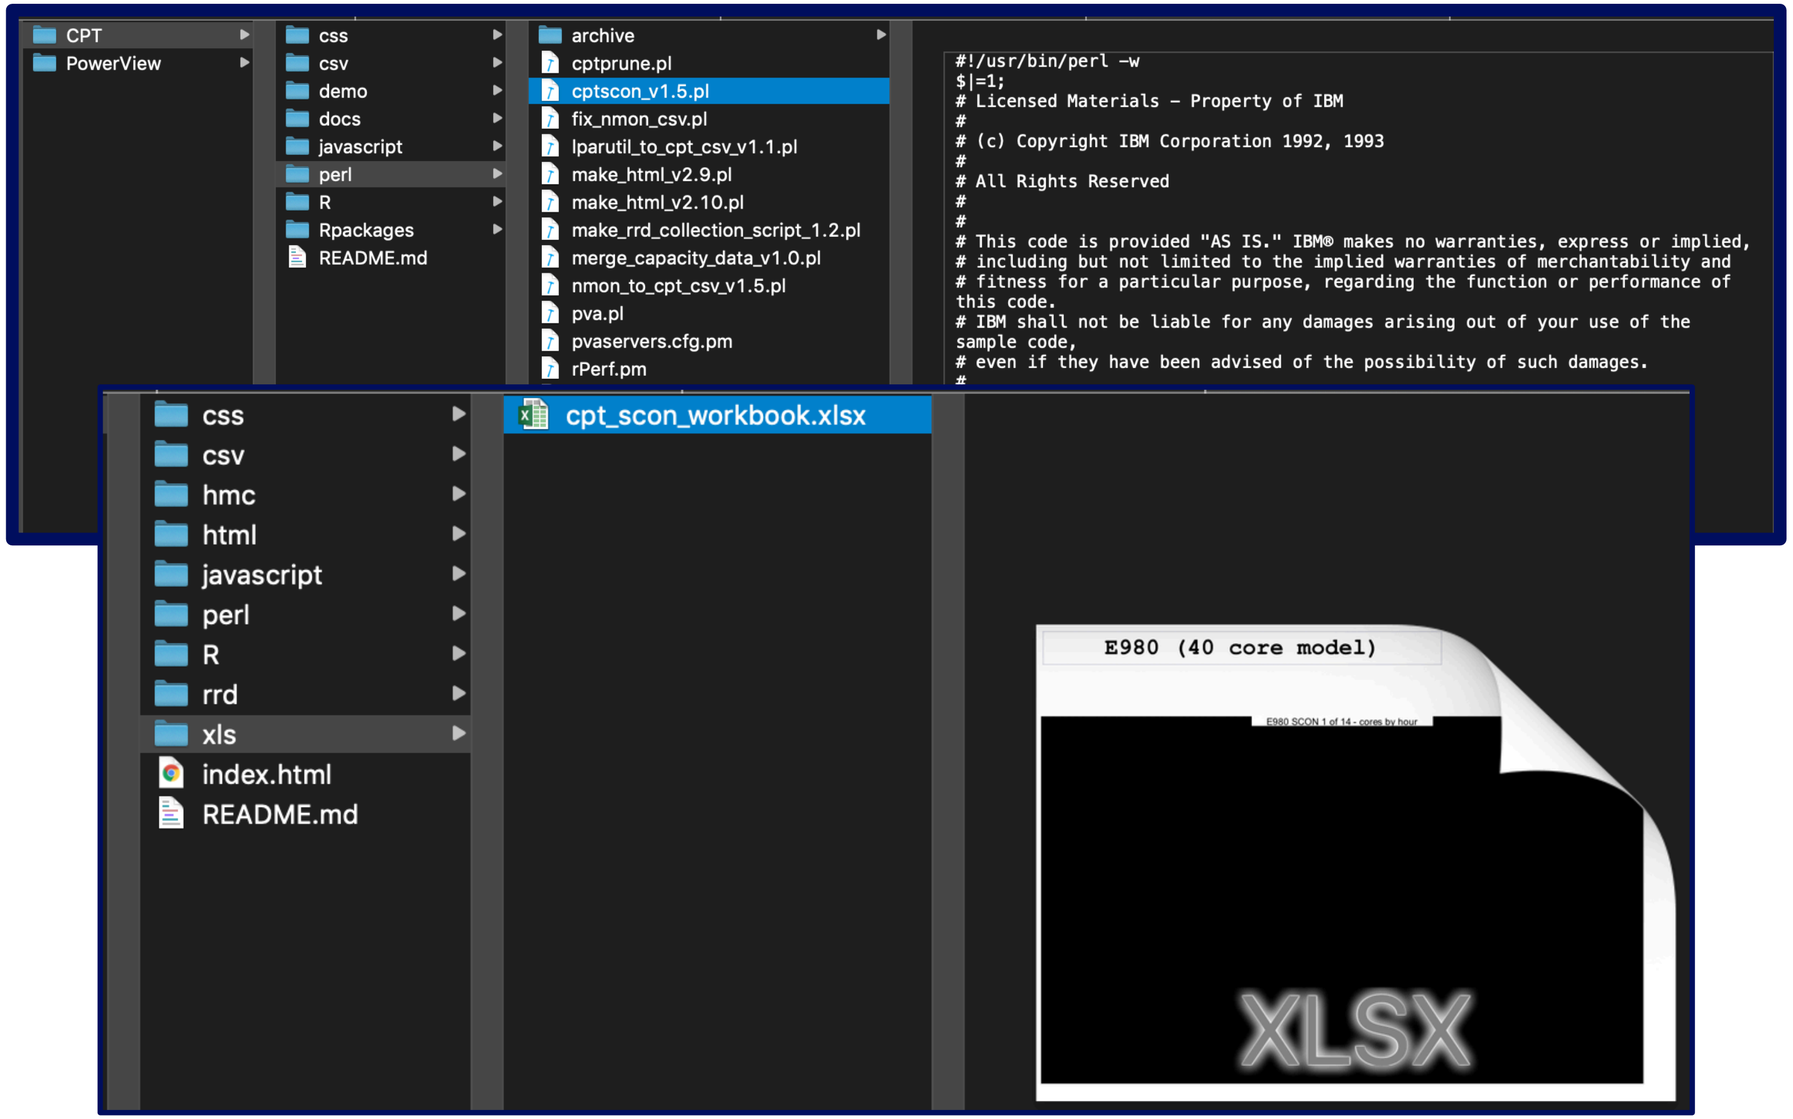

Capacity Planning Tool – SCONN Output

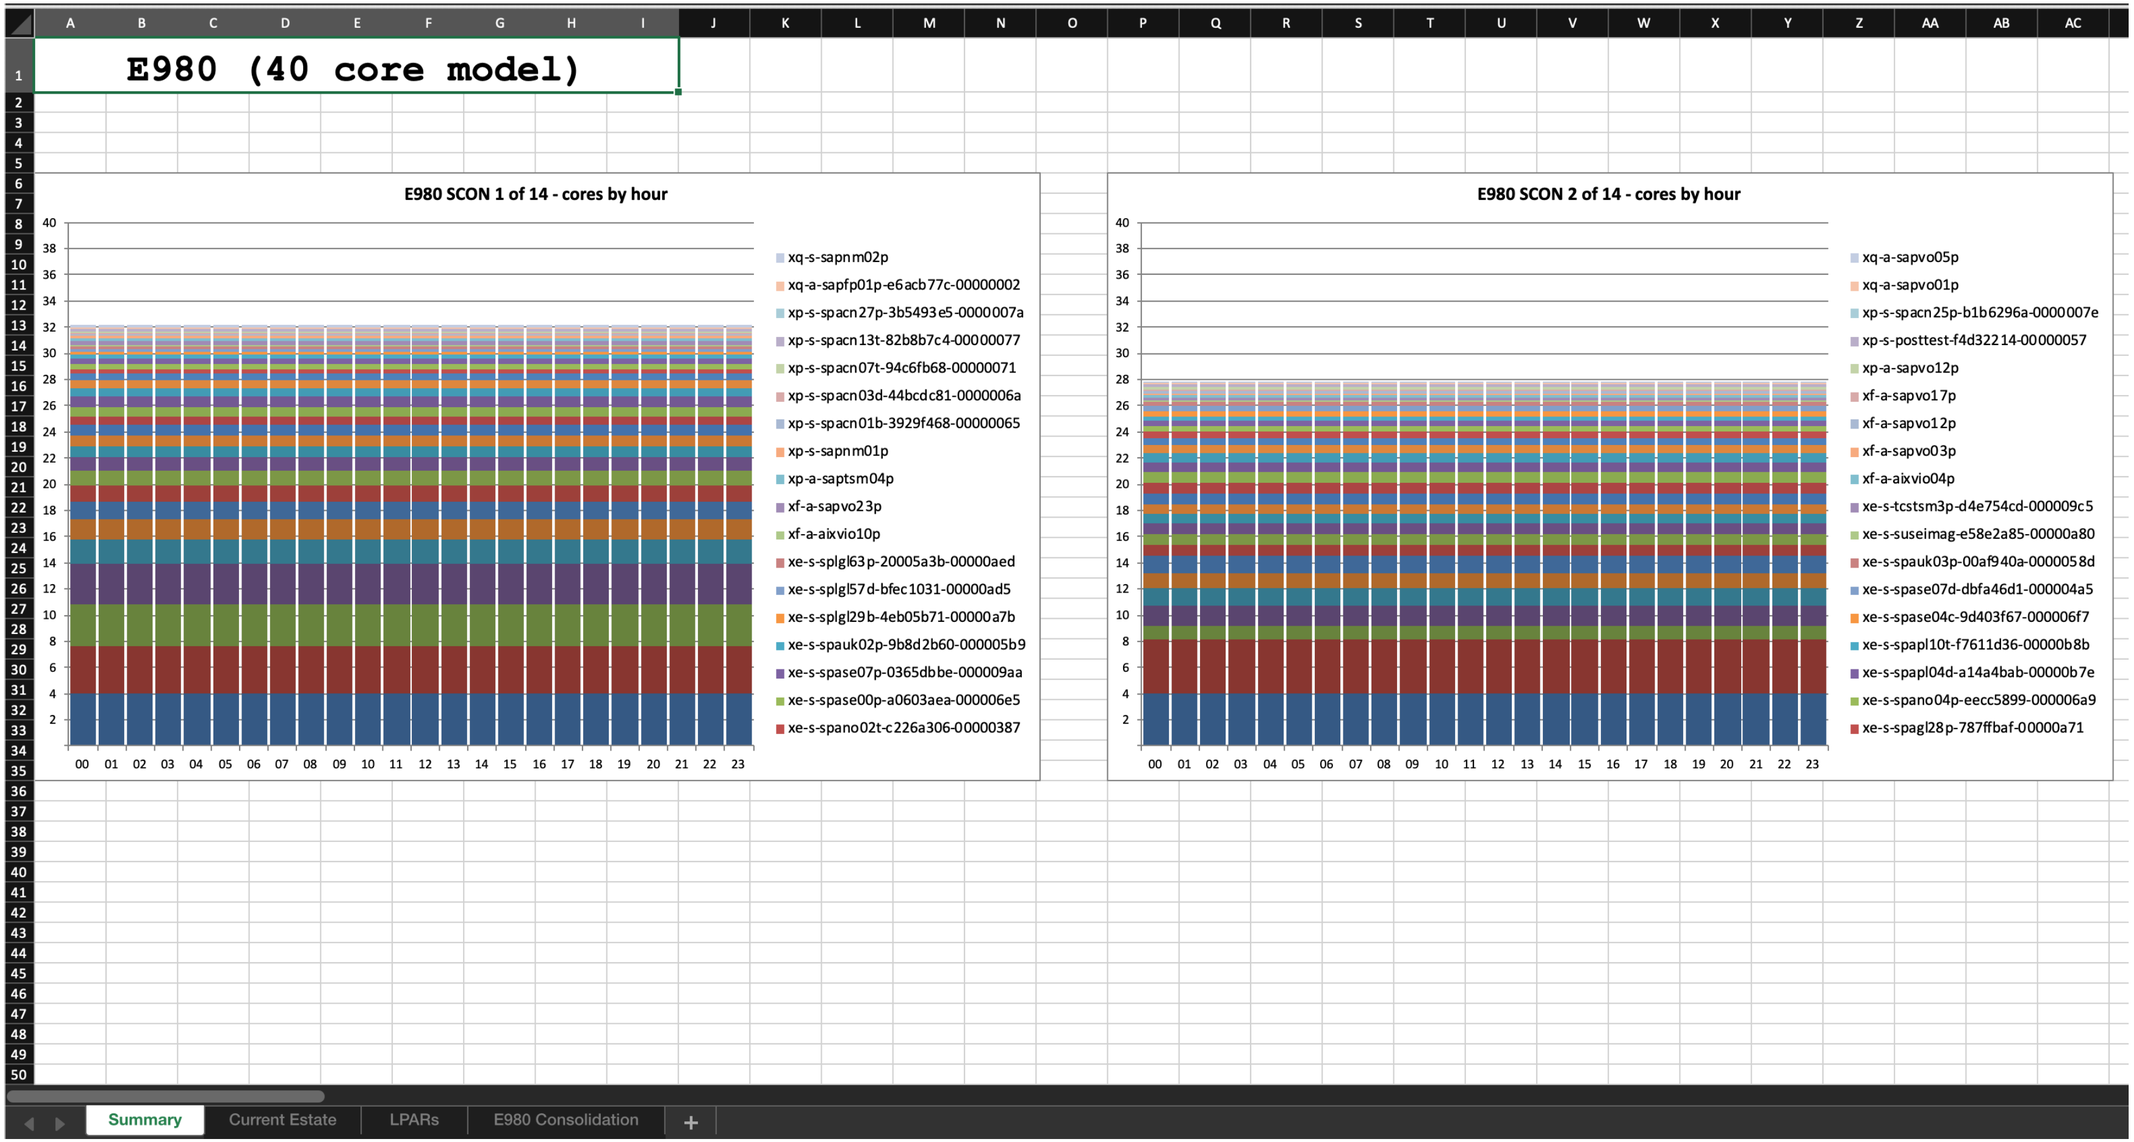

Capacity Planning Toolkit - Metered Capacity Modeling

CPT-MCM

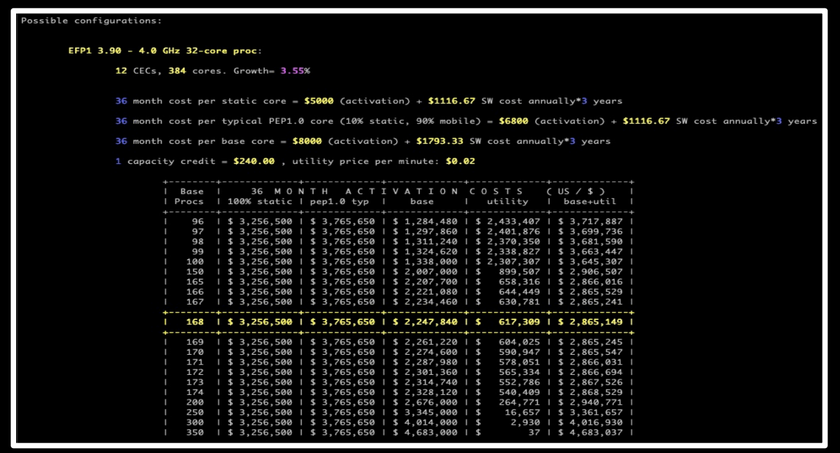

Possible configurations:

EFP1 3.90 - 4.0 GHz 32-core proc:

12 CECs, 384 cores. Growth= 3.55%

36 month cost per static core = $5000 (activation) + $1116.67 SW cost annually*3 years

36 month cost per typical PEP1.0 core (10% static, 90% mobile) = $6800 (activation) + $1116.67 SW cost annually*3 years

36 month cost per base core = $8000 (activation) + $1793.33 SW cost annually*3 years

1 capacity credit = $240.00 , utility price per minute: $0.02

Base Procs	36 MONTH ACTIVATION COSTS (US / $) 100% static	pep1.0 typ	base	utility	base+util
96	$ 3,256,500	$ 3,765,650	$ 1,284,480	$ 2,433,407	$ 3,717,887
97	$ 3,256,500	$ 3,765,650	$ 1,297,860	$ 2,401,876	$ 3,699,736
98	$ 3,256,500	$ 3,765,650	$ 1,311,240	$ 2,370,350	$ 3,681,590
99	$ 3,256,500	$ 3,765,650	$ 1,324,620	$ 2,338,827	$ 3,663,447
100	$ 3,256,500	$ 3,765,650	$ 1,338,000	$ 2,307,307	$ 3,645,307
150	$ 3,256,500	$ 3,765,650	$ 2,007,000	$ 899,507	$ 2,906,507
165	$ 3,256,500	$ 3,765,650	$ 2,207,700	$ 658,316	$ 2,866,016
166	$ 3,256,500	$ 3,765,650	$ 2,221,080	$ 644,449	$ 2,865,529
167	$ 3,256,500	$ 3,765,650	$ 2,234,460	$ 630,781	$ 2,865,241
168	$ 3,256,500	$ 3,765,650	$ 2,247,840	$ 617,309	$ 2,865,149
169	$ 3,256,500	$ 3,765,650	$ 2,261,220	$ 604,025	$ 2,865,245
170	$ 3,256,500	$ 3,765,650	$ 2,274,600	$ 590,947	$ 2,865,547
171	$ 3,256,500	$ 3,765,650	$ 2,287,980	$ 578,051	$ 2,866,031
172	$ 3,256,500	$ 3,765,650	$ 2,301,360	$ 565,334	$ 2,866,694
173	$ 3,256,500	$ 3,765,650	$ 2,314,740	$ 552,786	$ 2,867,526
174	$ 3,256,500	$ 3,765,650	$ 2,328,120	$ 540,409	$ 2,868,529
200	$ 3,256,500	$ 3,765,650	$ 2,676,000	$ 264,771	$ 2,940,771
250	$ 3,256,500	$ 3,765,650	$ 3,345,000	$ 16,657	$ 3,361,657
300	$ 3,256,500	$ 3,765,650	$ 4,014,000	$ 2,930	$ 4,016,930
350	$ 3,256,500	$ 3,765,650	$ 4,683,000	$ 37	$ 4,683,037

Designed to give customers insight into the potential benefits of the Power Enterprise Pools 2.0 offering.

Understand the benefits of PEP 2.0 before making the investment.

Use historical data from legacy systems and see rPerf scaled core utilization relative to E980 systems.

PEP 2.0 value calculated from historical CPU consumption data (one minute samples)

Model legacy workloads in latest technology using a Metered Capacity model

- Cost comparisons with full capacity investment and classic PEP models
- Optimize mix of base and metered capacity

Power Systems Private Cloud with Shared Utility Capacity

Cloud-like agility and economics with leadership business continuity and security

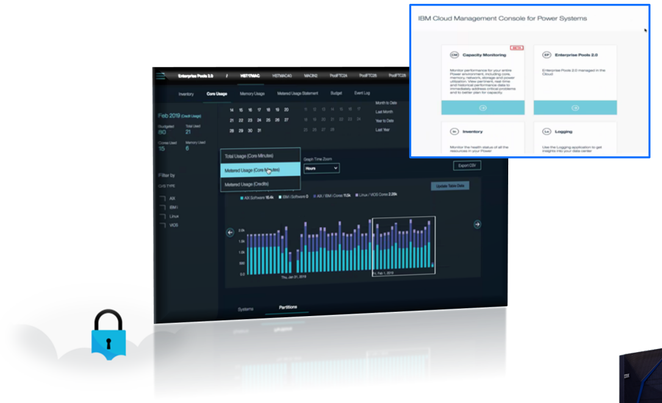

Power S922 & S924

Power E950

Power E980

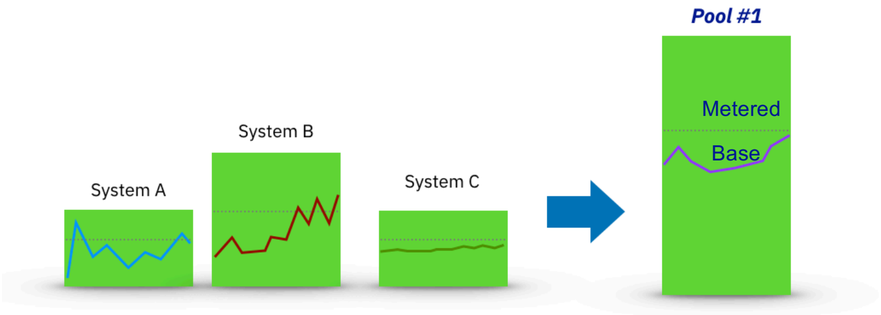

* One server machine type per pool. Multiple pools may be managed by a single instance of a Cloud Management Console

Expanded Shared Utility Capacity

- Deploy a Power Private Cloud infrastructure with Shared Utility Capacity across a collection of Power E980, or E950, or S924 and S922 systems*
- New, minimal system purchase/lease option as low as 1 core, 256GB active, with pay-per-use on balance of fully active capacity by the minute
- Industry-leading monitoring and metering via IBM Cloud Management Console with granular, real-time & historical views of consumption by resource by VM & system
- IBM Proactive Support
- Private Cloud Capacity Assessment & Implementation Services

Deploying Shared Utility Capacity

- Purchase servers with Base capacity
- Variable demand addressed by purchasing Capacity Credits for Metered capacity
- IBM Cloud Management Console with HMC automatically monitors and debits against Capacity Credits based on actual usage by the minute

Base and Metered Capacity

Processor activations

AIX and IBM i licenses

Memory activations
(E980 and E950 only)

Server Consolidation without Pooling or By-minute Metering

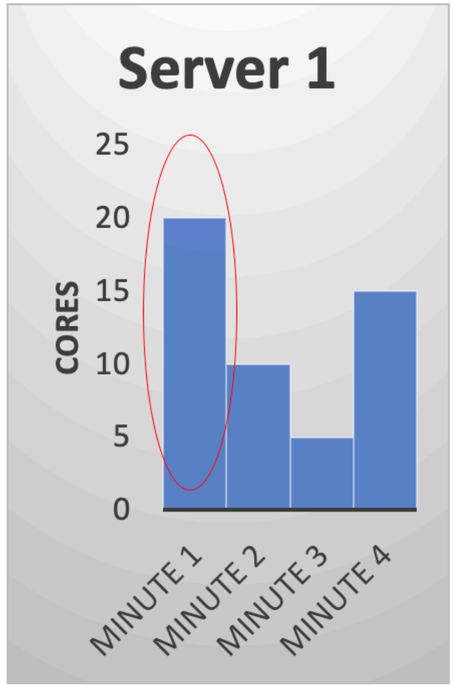

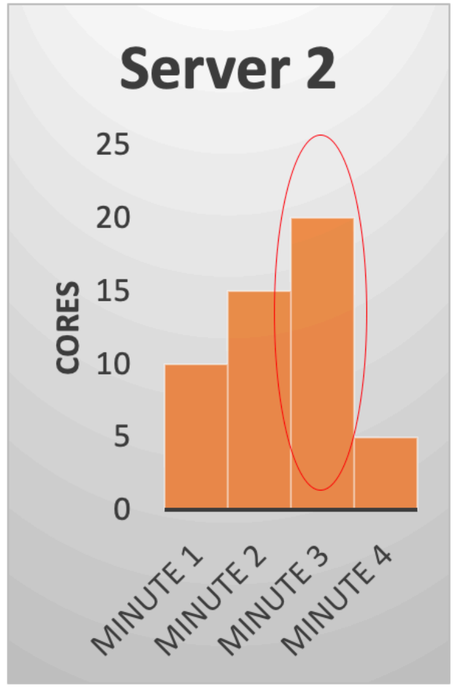

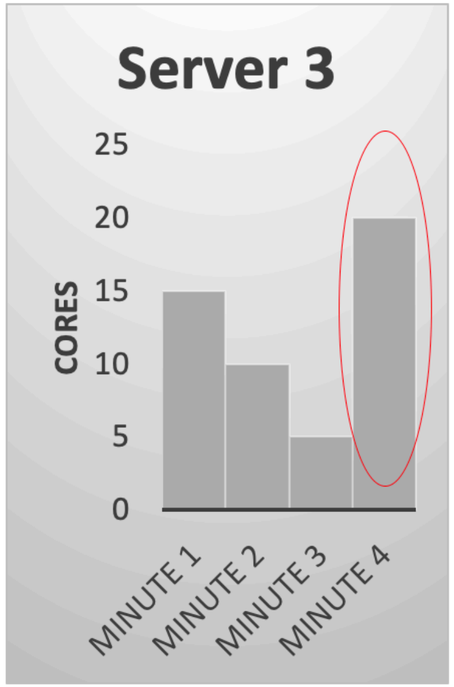

Consolidated system(s) must account for sum of all peaks (worst case assumption)

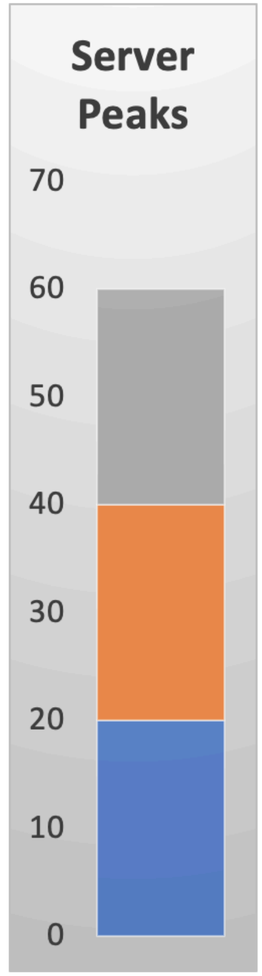

Server Consolidation *with* Pooling and By-minute Metering

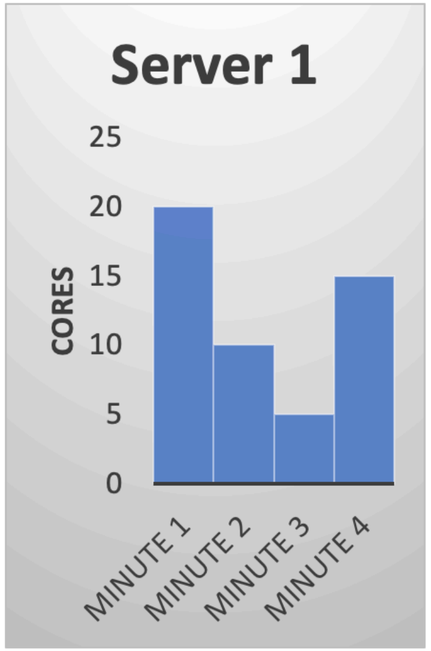

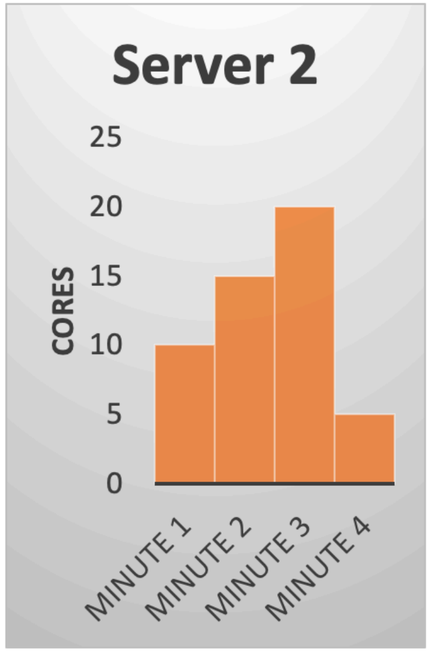

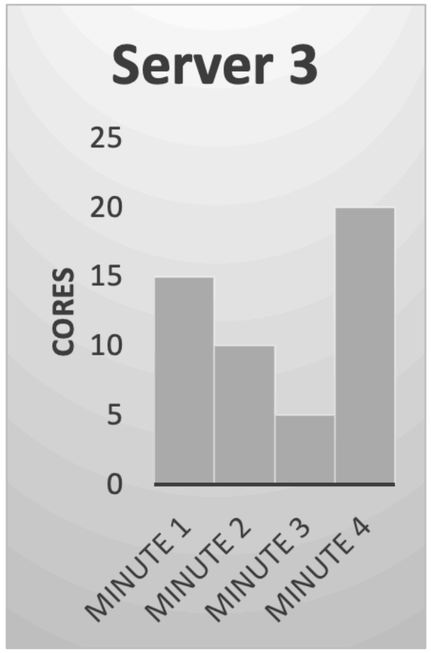

Pool peak is now based on aggregate per-minute usage

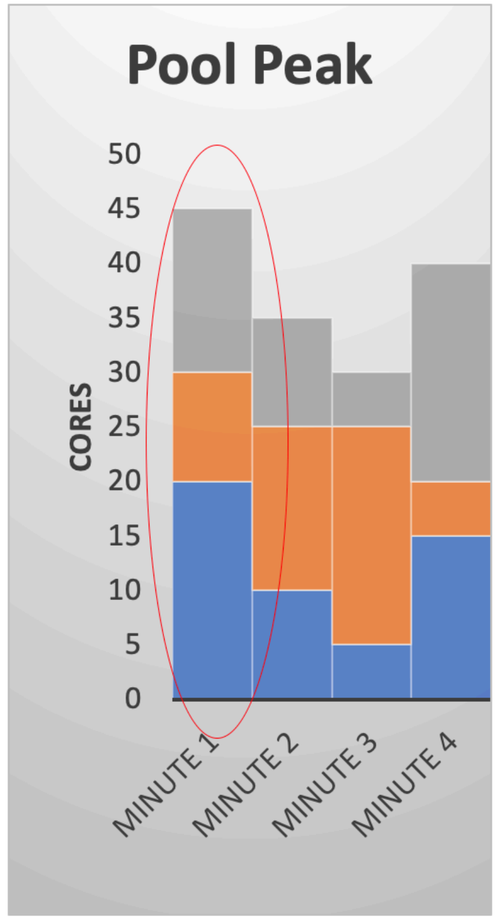

Processor Example - Pool has 1 system using more than its Base Processor Activations, but another system is idle, using less than its Base Processor resources at the same time, so 0 Metered resource usage is recorded

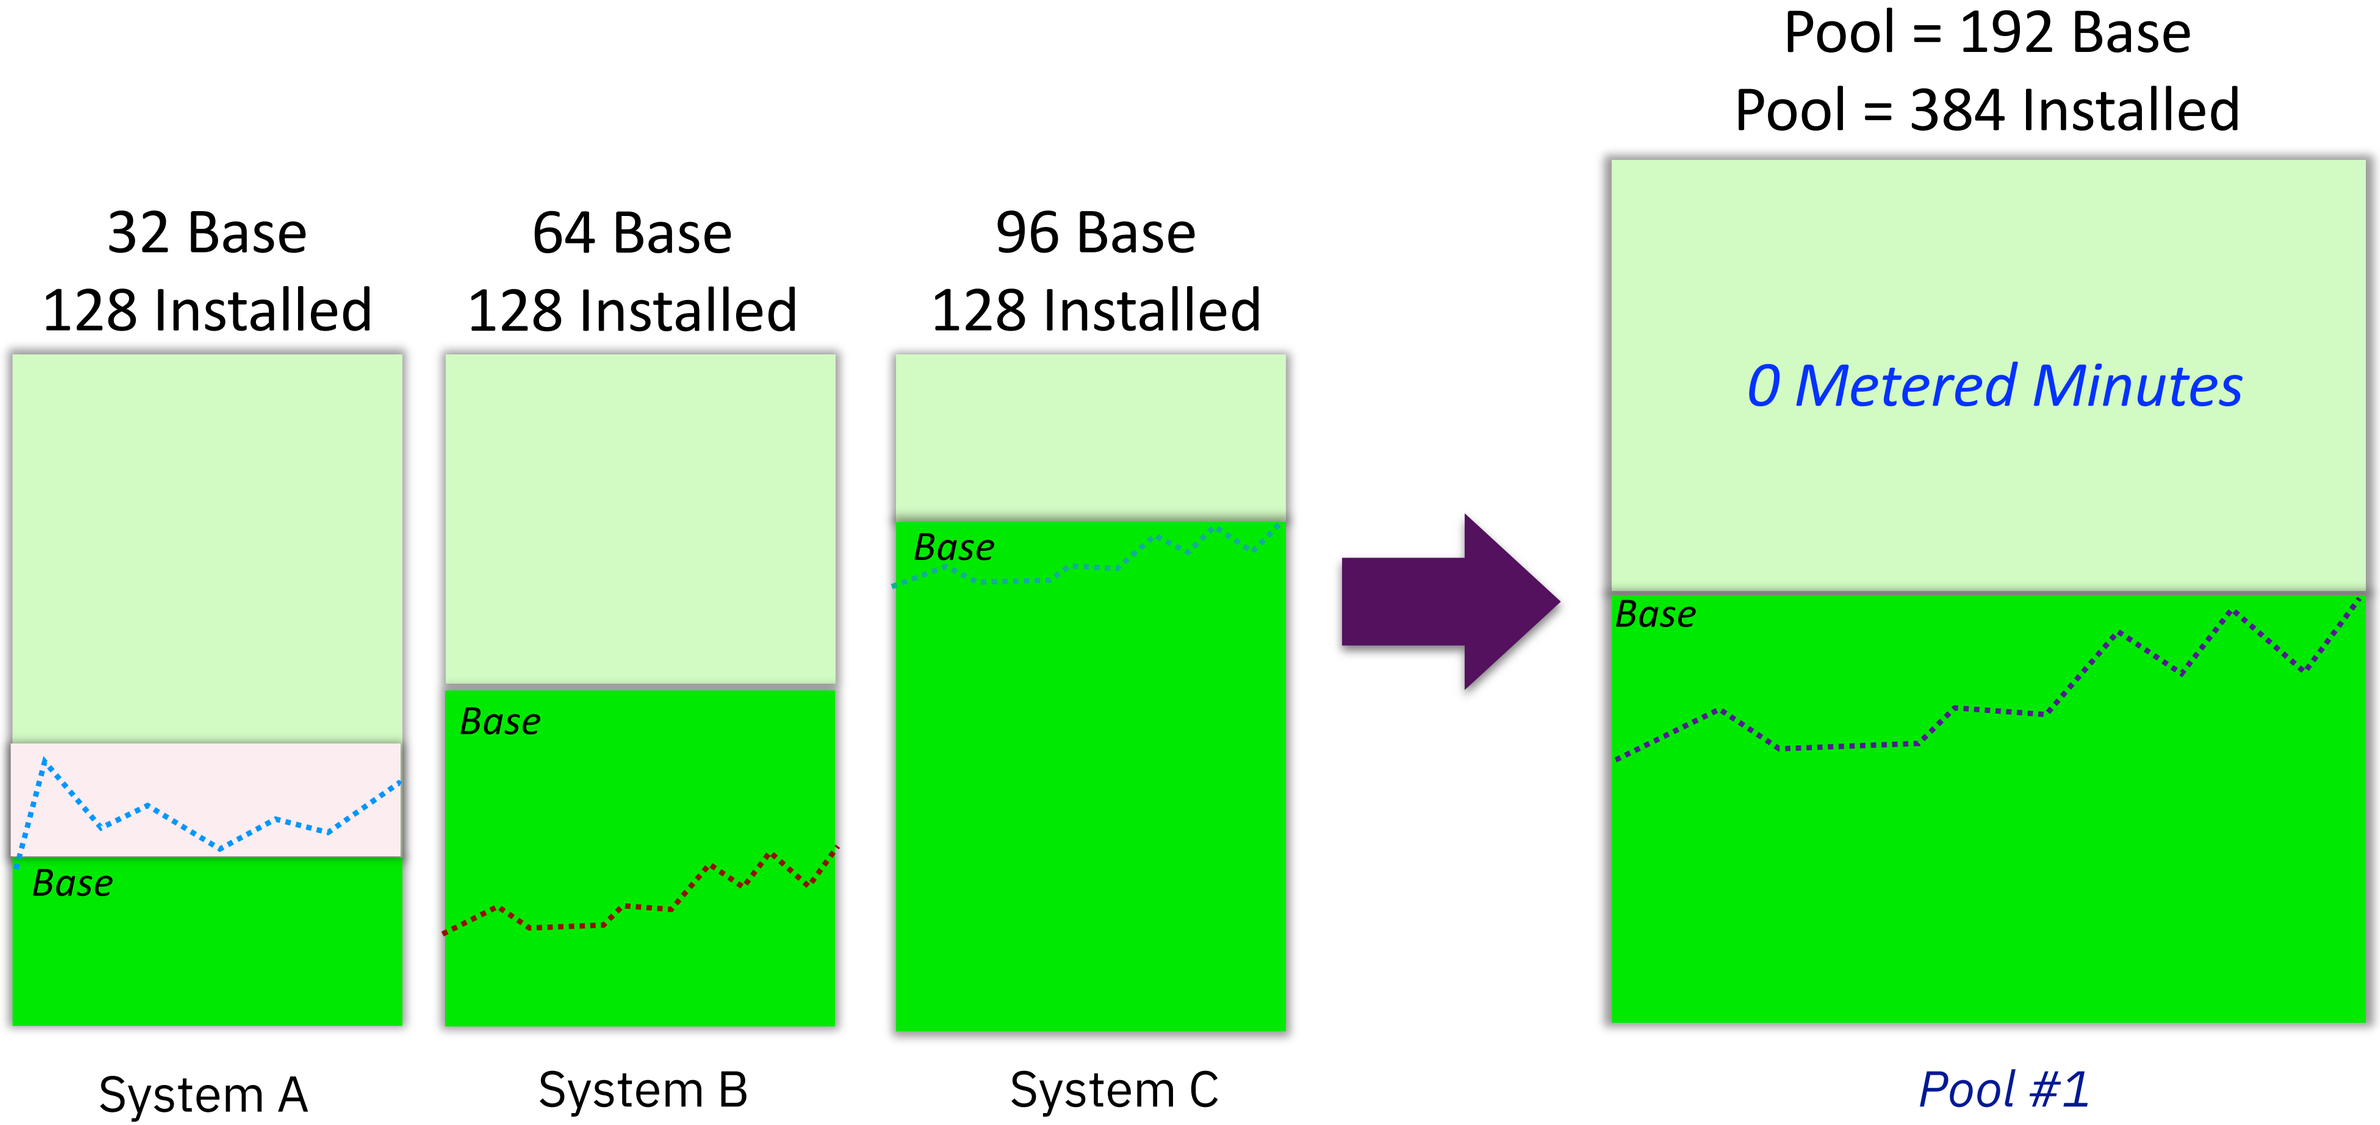

Pool’s base resources vs aggregated consumption

Pool = 192 Base
Pool = 384
Installed

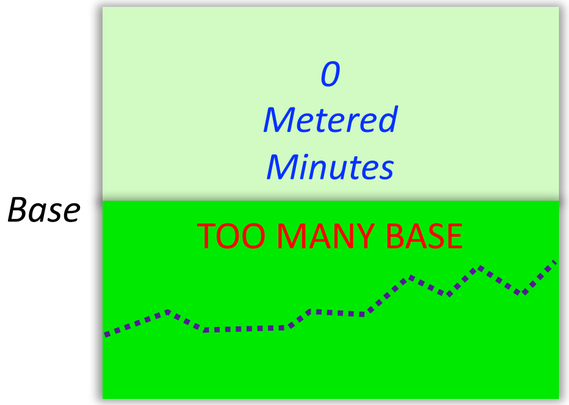

Pool #1

Unnecessary investment

Pool = 192 Base
Pool = 384
Installed

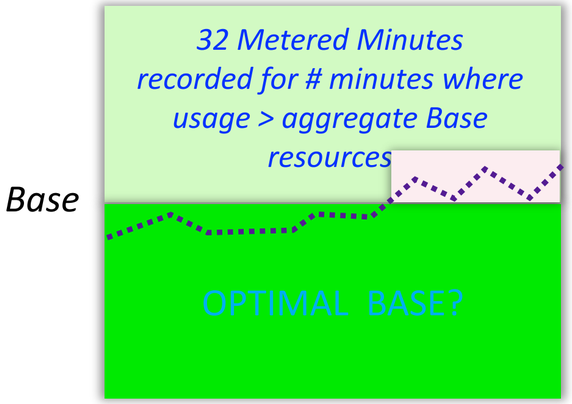

Pool #2

Ideal investment

Pool = 192 Base
Pool = 384
Installed

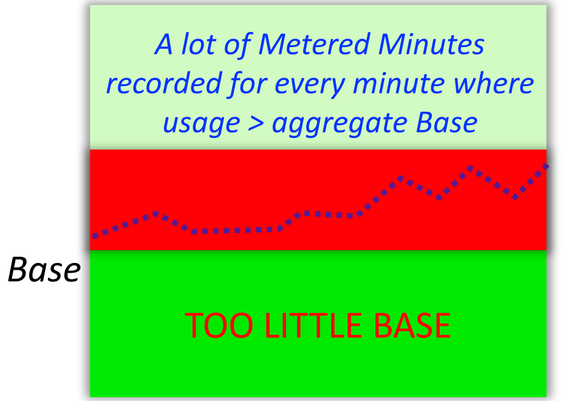

Pool #3

Potentially too expensive

CPT-MCM: Example Run

```
Step 1: Process server config file and retrieve rPerf info...

Step 2: Process server xml pool data - find start and stop times.................done.
            Start time: 1567959360 09/08/2019 16:16:00
              End time: 1570551240 10/08/2019 16:14:00
         Total Seconds: 2591880
         Total Minutes: 43198
           Total Hours: 719
            Total Days: 29

Step 3: Process server xml pool data - store physC data...
         Processing server FL1-9119-MHE-SN21C7867
         Processing server FL2-9119-MHE-SN21C7877
         Processing server FL3-9119-MHE-SN21C7887
         Processing server FL4-9119-MHE-SN21C7897
         Processing server FL5-9119-MHE-SN21C78A7
         Processing server FL6-9119-MHE-SN21C78B7
         Processing server GR1-9119-MME-SN2198FC7
         Processing server GR2-9119-MME-SN2198FD7
         Processing server GR3-9119-MME-SN2198FE7
         Processing server GR4-9119-MME-SN2198FF7
         Processing server GR5-9119-MME-SN2199007
         Processing server GR6-9119-MME-SN2199027
```

CPT-MCM: Example Run

```
Step 4: Calculate per minute usage...
                    FL1-9119-MHE-SN21C7867  (peak=47.89/36 time=09/21/2019 15:42:00)
                    FL2-9119-MHE-SN21C7877  (peak=45.81/25 time=09/21/2019 15:42:00)
                    FL3-9119-MHE-SN21C7887  (peak=56.13/37 time=09/09/2019 16:38:00)
                    FL4-9119-MHE-SN21C7897  (peak=48.00/50 time=10/05/2019 02:02:00)
                    FL5-9119-MHE-SN21C78A7  (peak=29.19/25 time=09/09/2019 16:36:00)
                    FL6-9119-MHE-SN21C78B7  (peak=67.53/59 time=10/03/2019 13:48:00)
                    GR1-9119-MME-SN2198FC7  (peak=31.95/28 time=09/21/2019 15:43:00)
                    GR2-9119-MME-SN2198FD7  (peak=24.10/23 time=09/11/2019 00:11:00)
                    GR3-9119-MME-SN2198FE7  (peak=27.19/24 time=09/21/2019 15:55:00)
                    GR4-9119-MME-SN2198FF7  (peak=34.28/27 time=10/05/2019 02:32:00)
                    GR5-9119-MME-SN2199007  (peak=43.07/38 time=09/25/2019 23:03:00)
                    GR6-9119-MME-SN2199027  (peak=40.23/36 time=10/05/2019 02:32:00)

      Worst case rPerf total  = 11,087 (this assumes no pooling or optimization via server consolidation)

              390 cores, EFP1 3.90 - 4.0 GHz 32-core proc (28.44 rPerf per core) $5,000/1 core activation for EFP1 $ 1,950,000
              404 cores, EFP2 3.70 - 3.9 GHz 40-core proc (27.45 rPerf per core) $5,000/1 core activation for EFP2 $ 2,020,000
              413 cores, EFP4 3.58 - 3.9 GHz 44-core proc (26.85 rPerf per core) $6,000/1 core activation for EFP4 $ 2,478,000
              396 cores, EFP3 3.55 - 3.9 GHz 48-core proc (28.01 rPerf per core) $6,000/1 core act for EFP3/EHC6   $ 2,376,000

      Peak minute rPerf total = 10,547 (10/05/2019 02:02:00) Best case is perfectly managed pool to this peak.

                                                                                                  50% static    10% static
                                                                                                  ----------    ----------
              371 cores, EFP1 3.90 - 4.0 GHz 32-core proc (28.44 rPerf per core) $7,000/1 core Mob Act EFP1/EFP5   $ 2,226,000 - $ 2,522,800
              385 cores, EFP2 3.70 - 3.9 GHz 40-core proc (27.45 rPerf per core) $7,000/1 core Mob Act EFP2/EFP6   $ 2,310,000 - $ 2,618,000
              393 cores, EFP4 3.58 - 3.9 GHz 44-core proc (26.85 rPerf per core) $7,000/1 core Mob Act EFP4/EFP8   $ 2,554,500 - $ 2,711,700
              377 cores, EFP3 3.55 - 3.9 GHz 48-core proc (28.01 rPerf per core) $7,000/1 core Mob Act EFP3/EFP7   $ 2,450,500 - $ 2,601,300
```

CPT-MCM: Pools Value Assessment Tool

SYSTEM SOFTWARE PER CORE PRICES				(based on 36 months owned)
Per-	BASE	STATIC	MOBILE	
Year	1793.33	1116.67	1116.67	
Minute	0.003410	0.002123	0.002123	(525960 minutes in a 365.25 day year)

CPT-MCM: Pools Value Assessment Tool

```
stgturgut:perl turgut$ ./pva.pl --dumppricing --country=DE


Pricing for country=DE:

CEC=EFP1 (3.90 - 4.0 GHz 32-core proc) €
              IFL ELBK (Linux proc act for EFP1/EFP5 ) €
             base EP90 (1-core Base Act Pools 2 EFP1 ) €
         baselinux EP96 (1-core BaseAct Linux EFP1    ) €
            mobile EFPE (1 core Mob Act EFP1/EFP5     ) €
            static EFPA (1 core activation for EFP1   ) €

CEC=EFP2 (3.70 - 3.9 GHz 40-core proc) €
              IFL ELBL (Linux proc act for EFP2/EFP6 ) €
             base EP91 (1-core Base Act Pools 2 EFP2 ) €
         baselinux EP97 (1-core BaseAct Linux EFP2    ) €
            mobile EFPF (1 core Mob Act EFP2/EFP6     ) €
            static EFPB (1 core activation for EFP2   ) €

CEC=EFP4 (3.58 - 3.9 GHz 44-core proc) €
              IFL ELBQ (Linux proc act for EFP4/EFP8 ) €
             base EP93 (1-core Base Act Pools 2 EFP4 ) €
         baselinux EP99 (1-core BaseAct Linux EFP4    ) €
            mobile EFPN (1 core Mob Act EFP4/EFP8     ) €
            static EFP9 (1 core activation for EFP4   ) €

CEC=EFP3 (3.55 - 3.9 GHz 48-core proc) €
              IFL ELBM (Linux proc act for EFP3/EFP7 ) €
             base EP92 (1-core Base Act Pools 2 EFP3 ) €
         baselinux EP98 (1-core BaseAct Linux EFP3    ) €
            mobile EFPG (1 core Mob Act EFP3/EFP7     ) €
            static EFPC (1 core act for EFP3/EHC6     ) €



PID=5765-AMT-0001 (AIX 7.2 Standard Edition Monthly Term Offering - Per Processor Core on Medium Server              ) €
    5771-AMT-2288 (AIX 7.2 Standard Edition Monthly Term Offering - 1 Year SWMA Per Processor Core on Medium Server  ) €

PID=5765-CBA-0001 (Enterprise Cloud Edition AIX - Per Processor Core on Medium Server                              ) €
    5771-CBA-2303 (Enterprise Cloud Edition AIX - 1 Year SWMA Per Processor Core on Medium Server                  ) €

PID=5765-CD3-0006 (AIX 7.2 ENTERPRISE EDITION V1 - Per Processor Core on Medium Server                            ) €
    5771-AEZ-1475 (AIX 7.2 ENTERPRISE EDITION V1 - 1 Year SWMA Per Processor Core on Medium Server                ) €

PID=5765-G98-0009 (IBM AIX Standard Edition V7 - Per Processor Core on Medium Server                               ) €
    5771-SWM-1510 (IBM AIX Standard Edition V7 - 1 Year SWMA Per Processor Core on Medium Server                   ) €
```

List prices are dynamically fetched by the tool.

Values can be overwritten with provided pricing information.

CPT-MCM: Pools Value Assessment Tool

Possible configurations:

EFP1 3.90 - 4.0 GHz 32-core proc:

12 CECs, **384** cores. Growth= **3.55%**

36 month cost per static core = **$5000** (activation) + **$1116.67** SW cost annually*3 years

36 month cost per typical PEP1.0 core (10% static, 90% mobile) = **$6800** (activation) + **$1116.67** SW cost annually*3 years

36 month cost per base core = **$8000** (activation) + **$1793.33** SW cost annually*3 years

1 capacity credit = **$240.00** , utility price per minute: **$0.02**

Base Procs	36 MONTH ACTIVATION COSTS (US / $) 100% static	pep1.0 typ	base	utility	base+util
96	$ 3,256,500	$ 3,765,650	$ 1,284,480	$ 2,433,407	$ 3,717,887
97	$ 3,256,500	$ 3,765,650	$ 1,297,860	$ 2,401,876	$ 3,699,736
98	$ 3,256,500	$ 3,765,650	$ 1,311,240	$ 2,370,350	$ 3,681,590
99	$ 3,256,500	$ 3,765,650	$ 1,324,620	$ 2,338,827	$ 3,663,447
100	$ 3,256,500	$ 3,765,650	$ 1,338,000	$ 2,307,307	$ 3,645,307
150	$ 3,256,500	$ 3,765,650	$ 2,007,000	$ 899,507	$ 2,906,507
165	$ 3,256,500	$ 3,765,650	$ 2,207,700	$ 658,316	$ 2,866,016
166	$ 3,256,500	$ 3,765,650	$ 2,221,080	$ 644,449	$ 2,865,529
167	$ 3,256,500	$ 3,765,650	$ 2,234,460	$ 630,781	$ 2,865,241
168	**$ 3,256,500**	**$ 3,765,650**	**$ 2,247,840**	**$ 617,309**	**$ 2,865,149**
169	$ 3,256,500	$ 3,765,650	$ 2,261,220	$ 604,025	$ 2,865,245
170	$ 3,256,500	$ 3,765,650	$ 2,274,600	$ 590,947	$ 2,865,547
171	$ 3,256,500	$ 3,765,650	$ 2,287,980	$ 578,051	$ 2,866,031
172	$ 3,256,500	$ 3,765,650	$ 2,301,360	$ 565,334	$ 2,866,694
173	$ 3,256,500	$ 3,765,650	$ 2,314,740	$ 552,786	$ 2,867,526
174	$ 3,256,500	$ 3,765,650	$ 2,328,120	$ 540,409	$ 2,868,529
200	$ 3,256,500	$ 3,765,650	$ 2,676,000	$ 264,771	$ 2,940,771
250	$ 3,256,500	$ 3,765,650	$ 3,345,000	$ 16,657	$ 3,361,657
300	$ 3,256,500	$ 3,765,650	$ 4,014,000	$ 2,930	$ 4,016,930
350	$ 3,256,500	$ 3,765,650	$ 4,683,000	$ 37	$ 4,683,037

Power Enterprise Pools 2.0 Redbook!

IBM

Redbooks IT Infrastructure Software Storage Training About Redbooks

IBM Redbooks > Power Systems >

IBM Power Systems Private Cloud with Shared Utility Capacity featuring Power Enterprise Pools 2.0

A draft IBM Redbooks publication

Updated 22 September 2020

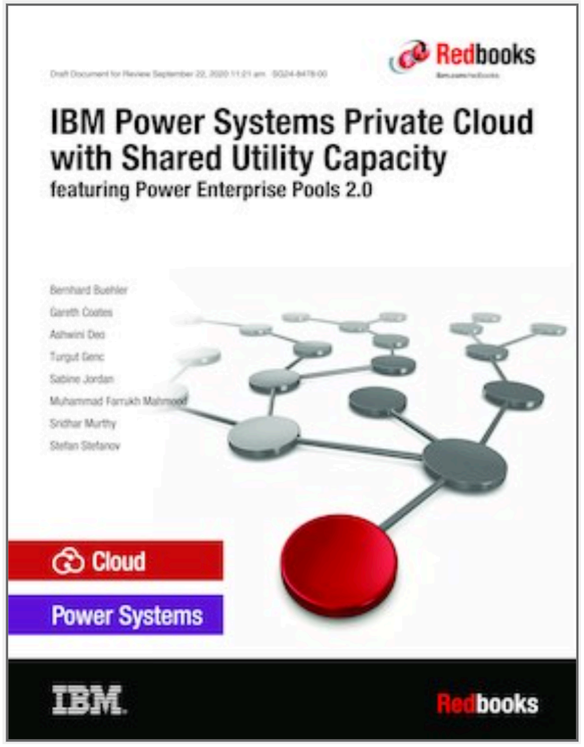

http://www.redbooks.ibm.com/abstracts/sg248478.html?Open

Contacts for obtaining Tools

- **Contacts and More Information**

  - **Lab Services UKI Opportunity Manager** : **John Markey** markejoh@ie.ibm.com
  - **Lab Services Europe Opportunity Manager** : **Beatrice Luquet** beatrice.luquet@fr.ibm.com
  - **Lab Services EMEA Technical Lead** : **Turgut Genc** TurgutGenc@uk.ibm.com

  - Other regions: please contact your local Lab Services opportunity manager http://ibm.biz/LabServicesOM
  - General Lab Services enquiries ibmsls@us.ibm.com

Demos

Questions

IBM Systems TechU

Featuring:

- Best practices from expert speakers and leading developers and consultants
- Technical sessions, hands-on labs and discounted certification testing
- Access to key technical decision makers and influencers
- Solution Centers showcasing the latest IBM and vendor products
- Dedicated time with clients, Business Partners, technical experts and developers

ibm.com/training/events

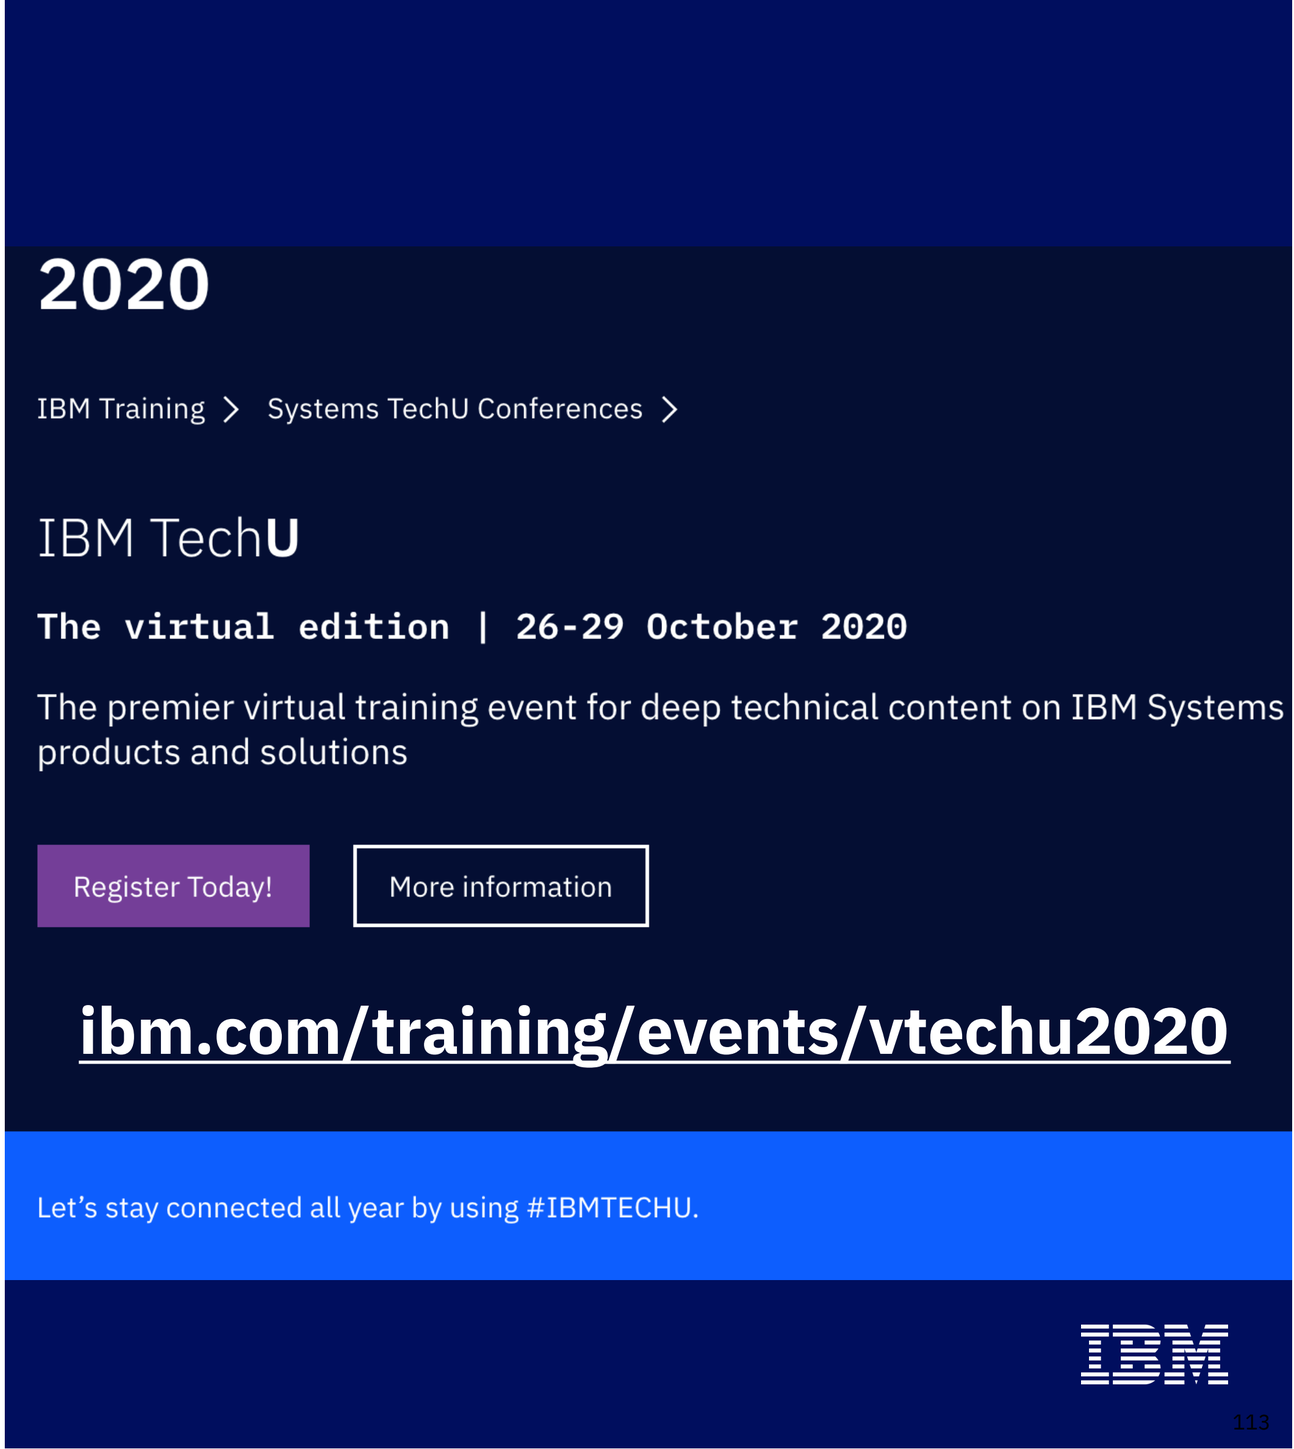

Lab Services and Business Partners

IBM Systems Lab Services:

- Helps BPs progress sales and accelerate adoption of IBM Systems
- Engages with and complements the skills of BPs
- Does not compete with or replace BP services practices
- Ensures BPs retain control of their client relationships

Working with IBM Systems Lab Services:

- Gives access to expertise of IBM consultants
- Gives access to skills transfer and shadowing
- Helps grow sales and extend business scope
- Helps reduce risk of working with new technologies

ibm.biz/LabServices

IBM Systems Lab Services

Infrastructure expertise to help you build the foundation of a smart enterprise.

IBM Systems Lab Services helps you deploy the building blocks of a next-generation IT infrastructure that empowers your business.

- Solve business challenges
- Gain new skills
- Apply best practices

Contact us today to see how Lab Services can help empower your business.

ibmsls@us.ibm.com

Resources

Visit:

Lab Services ibm.com website www.ibm.com/it-infrastructure/services/lab-services

Lab Services PartnerWorld website www.ibm.com/partnerworld/systems/services/lab-services

Lab Services Events website www.ibm.com/training/events

Lab Services developerWorks ibm.biz/LabServicesdW

IT Infrastructure blog channel www.ibm.com/blogs/systems/topics/services/ibm-systems-lab-services

Explore:

ibm.biz/LabServicesBrochure
ibm.biz/LabServicesPowerBrochure
ibm.biz/LabServicesStorageBrochure
ibm.biz/LabServiceszSystemsBrochure

Engage:

facebook.com/ibmsystems
twitter.com/IBMSystems
www.youtube.com/user/IBMSLST
linkedin.com/groups/8518569
#IBMLabServices #IBMTechU

Email:

ibmsls@us.ibm.com

Thank You