Common Practices & Optimization for Linux & SAP HANA

Torsten Wendland

(twendlan@de.ibm.com)

IBM Systems Lab Services Germany

Agenda

- Network Considerations
  - SAP HANA System Replication
  - Cluster (Pacemaker)
  - When should I think about SR-IOV?
  - General Topics
- Storage
  - Comparing VSCSI vs. NPIV with HWCCT
  - Multipath common issues
  - VIOS Tuning
  - Linux Tuning
  - Why Striping HANA DATA & LOG?
- Common Topics
  - SLES 12 SP1+ Tuned Common Issues
  - How to track SAP Note Changes?
  - Optimize Your Shell Code for OS Change
- IBM SLS Power To Cloud Services for SAP HANA on Power

Network Consideration

SAP HANA System Replication

- It is recommended to use at least a 10Gbit network
- Use a separate VLAN and physical network ports as your data network uses
  - Depending on Replicate and Operation Mode, HSR and Data might “steal” the throughput and latency from each other
  - E.g. Delta Data Shipping generates much more network traffic
  - E.g. SYNC mode requires low network response time, otherwise transaction throughput will be slowed down
  - In case of a full sync is needed between nodes, you will be happy about getting full throughput on your line and not sharing it with data
- HSR network can be a SEA, no problem with this
- Some customers sharing HSR and Backup network → this is fine as long as backup doesn’t conflict with huge amount of transactions at the same time → Topic Job Schedules

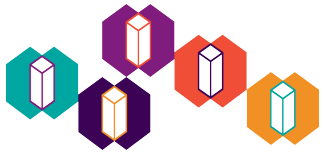

Cluster (Pacemaker)

- Pacemaker is the most used option as Cluster software for SAP HANA on Linux
- A lot more options are available like IBM PowerHA and other vendors with their pros & cons
- Pacemaker was never designed to be a 2-node cluster
- STONITH got introduced as a fencing mechanism in a 2-node cluster
- STONITH is not part of Pacemaker, but is used as fencing mechanism
  - It is handled as a (special) resource within the cluster
  - On Power, HMC and SAN volume are supported (more options are available, but usually not useable with Power)
  - Avoid using HMC as STONITH at much as possible → HMC SONITH is stateless and doesn’t work in case of a network split between DCs

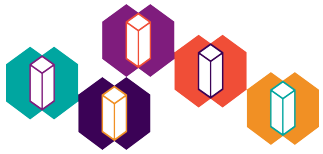

Cluster (Pacemaker)

- Pacemaker just uses network for communication and node/network down detection
- This requires proper planning to have at least two network “rings” within the cluster for proper node down/network down detection → Topic Split Brain Situation
- Highly recommended that each ring has a separate VLAN, and physical port(s)
  - If you already followed the recommendation for HANA System Replication, you already have your second network ring
- Physical ports of each ring (highly) should be attached to different switches/Fexes/Leafs

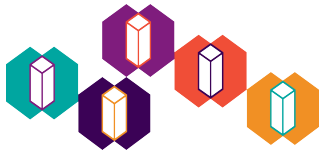

Cluster (Pacemaker)

- Pacemaker can only operate a 2-node SAP HANA design
  - You can have more nodes, but they cannot be part of the cluster
  - This applies for RHEL and SLES at this point in time

- Choose the right cluster software for your needs/requirements and not try to “bend” the cluster software, to fit in the architecture you have designed
- During the development of your architecture, review all cluster software options to verify that one of them can achieve what you have in mind

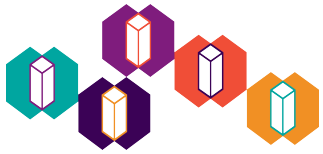

When Should I Think About SR-IOV?

- SR-IOV should be considered as an alternative beside SEA, specially for SAP LPARs with high amount of network transactions rates
- A classic application which benefits from SR-IOV, is ERP
- Evaluate the technology and it's different ways of implementation
  - e.g. SR-IOV VNIC is needed, if LPM is required
  - Don't forget about redundancy with SR-IOV and spare LPs for LPM

When Should I Think About SR-IOV?

- SAP ERP sizing for high network transaction rates.
- Packet size of 700 bytes assumed.
- Systems: Power8 E870 / Power7+ 770

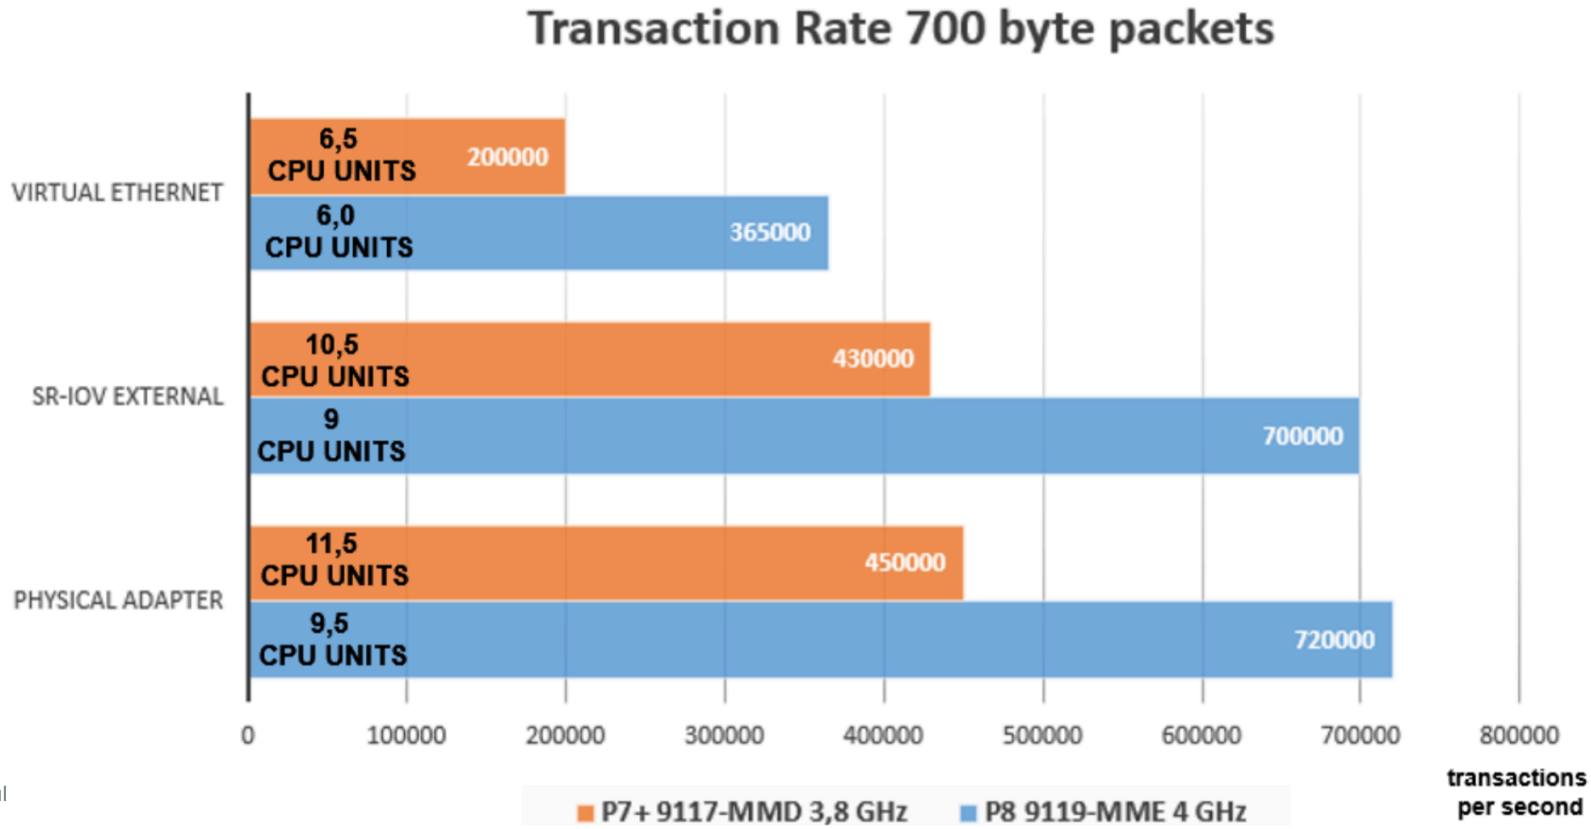

Source: Alexander Paul

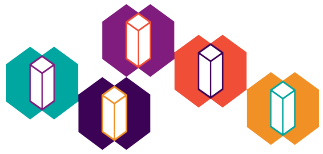

General Topics

- Don’t forget to configure large send & receive from VIOS up to the LPARs for your 10+ Gbit networks
- Increase the buffers of your SEAs for 10+ Gbit network → use the VIOS rules command as a good starting point
  - rules -o diff -d

Storage

Comparing VSCSI vs. NPIV with HWCCT

VSCSI attached LUNS with Active/Active on Adapters - 16Gbit
4 LUNs for each filesystem, Data striped, Log striped

Target KPIs

Volume	Block Sizes	Initial Write	Overwrite	Read	Latency
Log	4K	n.a.	30MB/s	n.a.	1000us
Log	16K	n.a.	120MB/s	n.a.	1000us
Log	1M	n.a.	250MB/s	250MB/s	n.a.
Data	4K	n.a.	n.a.	n.a.	n.a.
Data	16K	40MB/s	100MB/s	n.a.	n.a.
Data	64K	100MB/s	150MB/s	250MB/s	n.a.
Data	1M	150MB/s	200MB/s	300MB/s	n.a.
Data	16M	200MB/s	250MB/s	400MB/s	n.a.
Data	64M	200MB/s	250MB/s	400MB/s	n.a.

Results

Volume	Block Sizes	Initial Write	Overwrite	Read	Latency
Log	4K	7MB/s	245MB/s	328MB/s	**318us**
Log	16K	39MB/s	**771MB/s**	802MB/s	**440us**
Log	1M	410MB/s	3846MB/s	3775MB/s	1304us
Data	4K	42MB/s	245MB/s	215MB/s	485us
Data	16K	156MB/s	777MB/s	364MB/s	506us
Data	64K	1048MB/s	444MB/s	1776MB/s	508us
Data	1M	3651MB/s	2725MB/s	2220MB/s	1622us
Data	16M	3716MB/s	2546MB/s	2755MB/s	10586us
Data	64M	2690MB/s	3877MB/s	3878MB/s	20000us

Comparing VSCSI vs. NPIV with HWCCT

NPIV attached LUNS with 8 Paths - 16Gbit
4 LUNs for each filesystem, Data striped, Log striped

Target KPIs

Volume	Block Sizes	Initial Write	Overwrite	Read	Latency
Log	4K	n.a.	30MB/s	n.a.	1000us
Log	16K	n.a.	120MB/s	n.a.	1000us
Log	1M	n.a.	250MB/s	250MB/s	n.a.
Data	4K	n.a.	n.a.	n.a.	n.a.
Data	16K	40MB/s	100MB/s	n.a.	n.a.
Data	64K	100MB/s	150MB/s	250MB/s	n.a.
Data	1M	150MB/s	200MB/s	300MB/s	n.a.
Data	16M	200MB/s	250MB/s	400MB/s	n.a.
Data	64M	200MB/s	250MB/s	400MB/s	n.a.

Results

Volume	Block Sizes	Initial Write	Overwrite	Read	Latency
Log	4K	12MB/s	275MB/s	331MB/s	**304us**
Log	16K	44MB/s	**787MB/s**	901MB/s	**344us**
Log	1M	713MB/s	3032MB/s	3087MB/s	1339us
Data	4K	63MB/s	280MB/s	289MB/s	333us
Data	16K	303MB/s	1031MB/s	1015MB/s	373us
Data	64K	1008MB/s	2817MB/s	3061MB/s	503us
Data	1M	3284MB/s	3474MB/s	3332MB/s	1669us
Data	16M	3009MB/s	3109MB/s	3443MB/s	7794us
Data	64M	2865MB/s	2921MB/s	3370MB/s	24736us

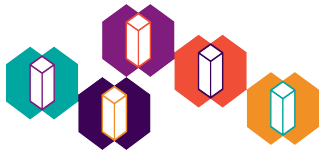

Comparing VSCSI vs. NPIV with HWCCT

- We would have expected a significant better performance with NPIV but the different tests showed that there is no difference in I/O performance between VSCSI vs NPIV attached volumes

- Due to the fact that Linux uses the VSCSI adapter in Active/Active mode, the sum of the available paths on the physical adapters is the same as you would attach the volumes by NPIV with 8 paths
  - The VSCSI test was run with 4 path per volume

- The possibility to use VSCSI in active/active mode on Linux is different then in AIX (only active/failover is supported)

- We got the confirmation of the developer, that the VSCSI is always running in active/active by default on Linux

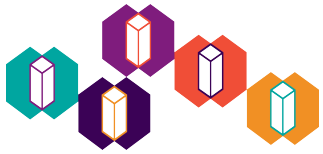

Multipath Common Issues

- Very often the multipathd is not configured and runs with the defaults, which are not always the recommended settings for most of the storage subsystems
- Consult the storage manufacturers website/support to get the latest Linux multipath settings
  - Sometime they can differ between RHEL and SLES and even within the same distributors offered versions
  - IBM Storwize products might have different settings depending on the storage firmware version installed (take care, this might also be true for some other manufacturers)
- Often customers are using user_frirendly_names = yes, but do not run an mkinitrd after any change (add volume/change volume name/delete volume)
  - This can lead to mixed names after the boot, because the multipath.conf and bindings file are included in the initrd

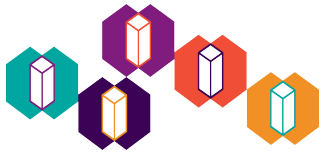

VIOS Tuning

- Still some customer run their fiber channel adapters with the defaults
- This limits the performance and capability of the new adapters, specially the 16Gbit ones
- Important parameters to change are
  - num_cmd_elems on the fcsX device
  - dyntrk on the fscsiX device → this change has nothing to do with performance, but with path reliability
  - max_transfer_rate on the fcsX device → specially for newer and more capable storage subsystems, but needs to be checked what is allowed/recommended from the storage provider
- An alternate way to check for common practices for the different devices models, use the VIOS rules command ☺
  - rules -o diff -d

Linux Tuning

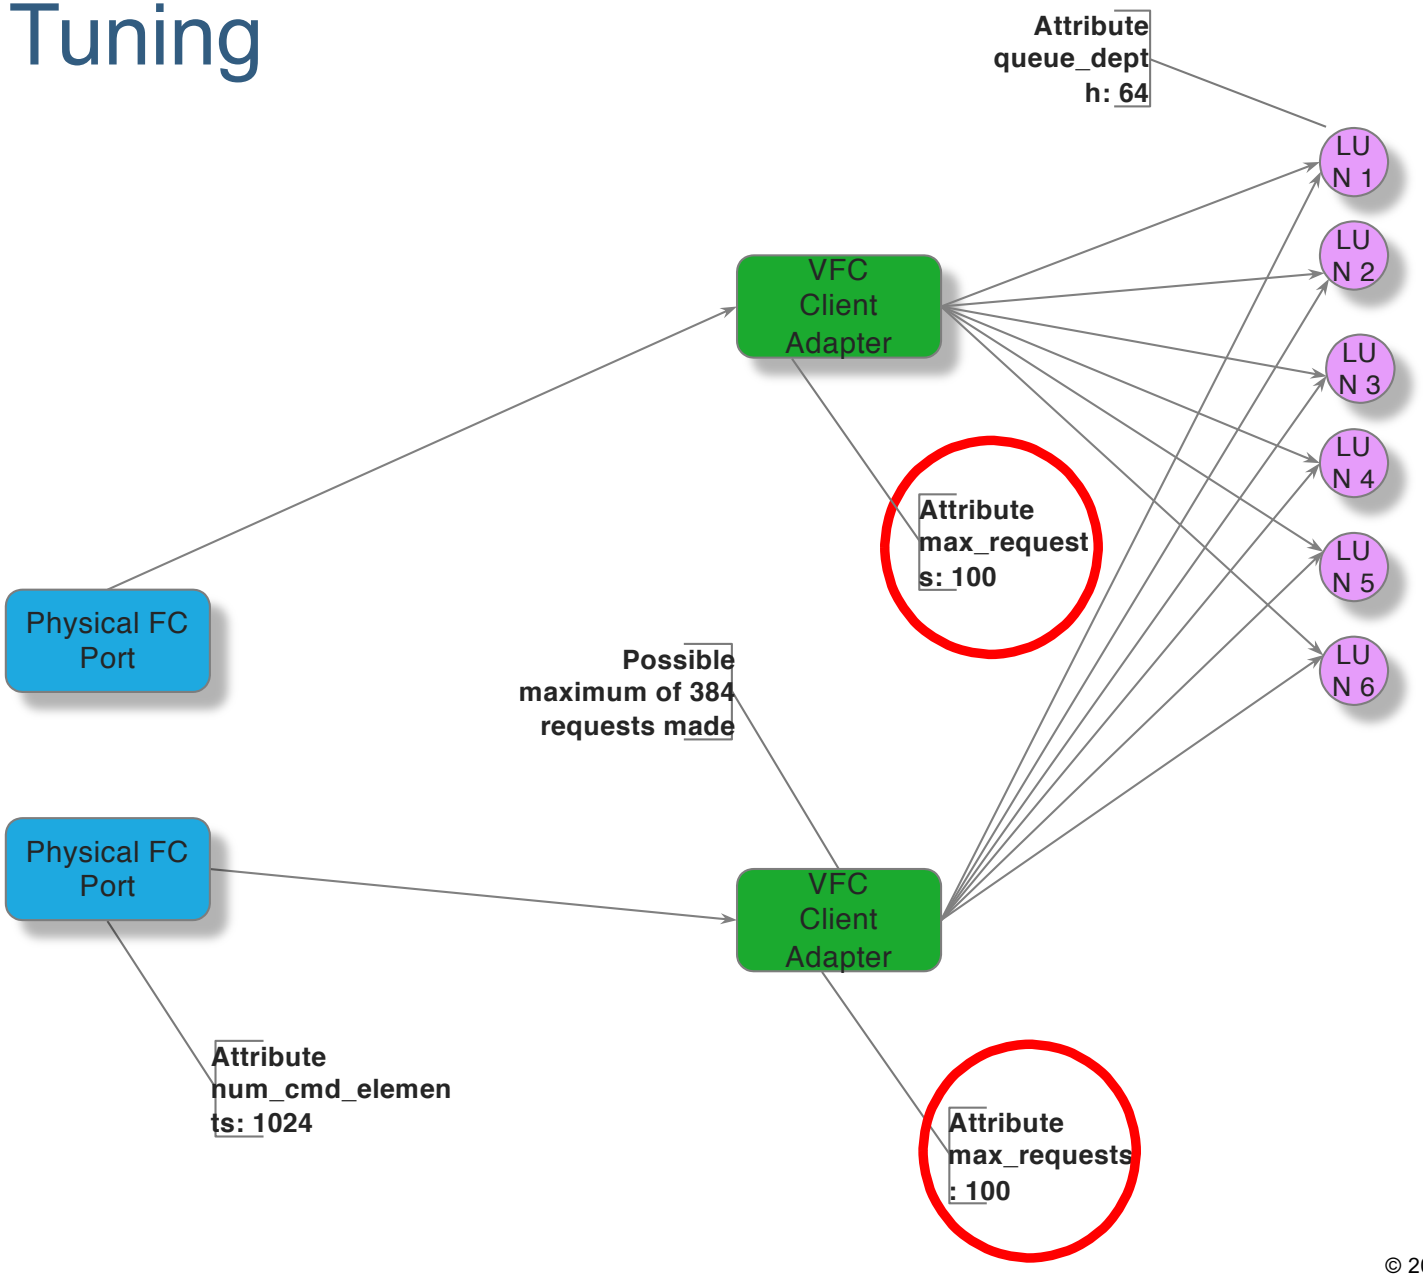

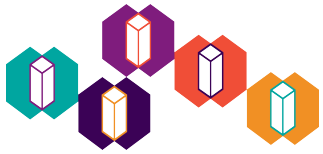

Linux Tuning

- Also Linux offers some tuning (a little bit like AIX)
- For NPIV the disk queue_depth can be increased, specially for 16Gbit SAN
  - Be careful with VSCSI, here the queue_depth should be left at the default, because VSCSI adapters have different limitations of how many SCSI commands / I/O requests it can accept
- Also for NPIV the queue size of the VFC adapter should be increased
- But keep in mind about overloading the whole stack (see page before)

Linux Tuning

- The VFC adapter queue can be increased by a kernel module parameter during the boot

 # vi /etc/default/grub

```
GRUB_CMDLINE_LINUX_DEFAULT="splash=silent quiet showopts elevator=noop numa_balancing=disable
ibmvfc.max_requests=200 transparent_hugepage=never crashkernel=256M-2G:64M,2G-:128M"
# Run the following 2 commands to make the change active for the next reboot (reboot is required to
activate it afterwards)
# grub2-mkconfig -o /boot/grub2/grub.cfg
# mkinitrd
```

reboot
cat /sys/module/ibmvfc/parameters/max_requests
200

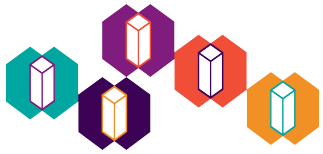

Linux Tuning

- The disk queue_depth can be changed with UDEV and needs some “rules coding”
- With SLES 12+ the read_ahead_kb should be changed by tuned, but I observed that it not always happens → apply the change by UDEV to ensure that it is always set
- This is an example for IBM V7K:

```
linux:/etc/udev/rules.d # vi /etc/udev/rules.d/99-ibm-2145.rules

# Set SCSI command timeout to 120s (default == 30 or 60) for IBM 2145 devices
SUBSYSTEM=="block", ACTION=="add", ENV{ID_VENDOR}=="IBM",ENV{ID_MODEL}=="2145", RUN+="/bin/sh -c 'echo 120
>/sys/block/%k/device/timeout'”

ACTION=="add|change",SUBSYSTEM=="block", ENV{ID_VENDOR}=="IBM",ENV{ID_MODEL}=="2145”,
ATTR{queue/read_ahead_kb}=”4096",ATTR{device/queue_depth}="64”

linux:/etc/udev/rules.d # udevadm control –reload-rules
linux:/etc/udev/rules.d # udevadm trigger

linux:/etc/udev/rules.d # for i in $(find -H /sys/block/sd* -name queue_depth); do cat $i; done
linux:/etc/udev/rules.d # for i in $(find -H /sys/block/sd*/queue -name read_ahead_kb); do cat $i; done
```

Linux Tuning

linux:/etc/udev/rules.d # udevadm info --query all --path /sys/block/sdb --attribute-walk

looking at device '/devices/vio/3000000a/host2/rport-2:0-1/target2:0:1/2:0:1:0/block/sdb':
KERNEL=="sdb"
SUBSYSTEM=="block"
DRIVER==""
<SNIP>
ATTRS{model}=="Invista "
ATTRS{queue_depth}=="32"
ATTRS{queue_ramp_up_period}=="120000"
ATTRS{queue_type}=="simple"
ATTRS{rev}=="6010" ATTRS{scsi_level}=="5"
ATTRS{state}=="running"
ATTRS{timeout}=="30"
ATTRS{type}=="0"
ATTRS{vendor}=="EMC "
<SNIP>

Why Stripping HANA DATA & LOG?

- Each volume has a single queue, which defines how many I/O requests in can take, until any further I/O gets delayed/blocked
- I/O requests in queues are handled sequentially (FIFO – First In, First Out)
- Having a filesystem placed on a single volume, only one queue will be available
- Placing a filesystem across more then one volume (without explicit striping) offers more queues, but still I/O requests are not balanced and still leads to hot-spots on single volumes
- Striping takes care of spreading the data as equally as possible across all volumes
  - This leads in a better utilization of all available volume queues and will result in better throughput (R/W) and latency due to parallelism of the I/O

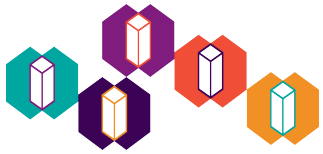

Why Stripping HANA DATA & LOG?

VSCSI attached LUNS with Active/Active on Adapters - 16Gbit 4 LUNs for each filesystem, Data striped, Log not striped

Target KPIs

Volume	Block Sizes	Initial Write	Overwrite	Read	Latency
Log	4K	n.a.	30MB/s	n.a.	1000us
Log	16K	n.a.	120MB/s	n.a.	1000us
Log	1M	n.a.	250MB/s	250MB/s	n.a.
Data	4K	n.a.	n.a.	n.a.	n.a.
Data	16K	40MB/s	100MB/s	n.a.	n.a.
Data	64K	100MB/s	150MB/s	250MB/s	n.a.
Data	1M	150MB/s	200MB/s	300MB/s	n.a.
Data	16M	200MB/s	250MB/s	400MB/s	n.a.
Data	64M	200MB/s	250MB/s	400MB/s	n.a.

Results

Volume	Block Sizes	Initial Write	Overwrite	Read	Latency
Log	4K	7MB/s	191MB/s	215MB/s	**513us**
Log	16K	26MB/s	**211MB/s**	205MB/s	**573us**
Log	1M	172MB/s	3091MB/s	3668MB/s	5592us
Data	4K	58MB/s	238MB/s	259MB/s	533us
Data	16K	233MB/s	879MB/s	583MB/s	606us
Data	64K	915MB/s	2230MB/s	2304MB/s	799us
Data	1M	2933MB/s	3428MB/s	3600MB/s	1935us
Data	16M	2107MB/s	2945MB/s	3714MB/s	10886us
Data	64M	2792MB/s	2945MB/s	3742MB/s	26559us

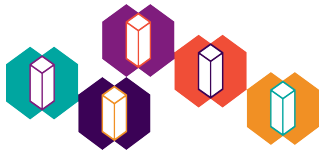

Why Stripping HANA DATA & LOG?

VSCSI attached LUNS with Active/Active on Adapters - 16Gbit
4 LUNs for each filesystem, Data striped, Log striped

Target KPIs

Volume	Block Sizes	Initial Write	Overwrite	Read	Latency
Log	4K	n.a.	30MB/s	n.a.	1000us
Log	16K	n.a.	120MB/s	n.a.	1000us
Log	1M	n.a.	250MB/s	250MB/s	n.a.
Data	4K	n.a.	n.a.	n.a.	n.a.
Data	16K	40MB/s	100MB/s	n.a.	n.a.
Data	64K	100MB/s	150MB/s	250MB/s	n.a.
Data	1M	150MB/s	200MB/s	300MB/s	n.a.
Data	16M	200MB/s	250MB/s	400MB/s	n.a.
Data	64M	200MB/s	250MB/s	400MB/s	n.a.

Results

Volume	Block Sizes	Initial Write	Overwrite	Read	Latency
Log	4K	7MB/s	245MB/s	328MB/s	**318us**
Log	16K	39MB/s	**771MB/s**	802MB/s	**440us**
Log	1M	410MB/s	3846MB/s	3775MB/s	1304us
Data	4K	42MB/s	245MB/s	215MB/s	485us
Data	16K	156MB/s	777MB/s	364MB/s	506us
Data	64K	1048MB/s	444MB/s	1776MB/s	508us
Data	1M	3651MB/s	2725MB/s	2220MB/s	1622us
Data	16M	3716MB/s	2546MB/s	2755MB/s	10586us
Data	64M	2690MB/s	3877MB/s	3878MB/s	20000us

Why Stripping HANA DATA & LOG?

- Striping the HANA LOG filesystem
  - Improves latency → transactions will benefit from it
  - Improves Overwrite Speed for HANA → transactions will benefit from it
  - Improves read throughput → crash recovery and HANA System Replication with log-shipping will benefit from it
- You already stipe the HANA DATA (recommendation from SAP), so striping HANA LOG too will not lead into much more handling costs compared to the improvements

Common Topics

SLES 12 SP1+ Tuned Common Issues

- A wide spread common issue has started with SLES 12 SP1+
- SLES 11 had the sapconf tool, to prepare some required/recommended settings
- SLES 12 has the tuned with profiles for simplification of setting tuning parameters
  - Very often tuned runs with the sap-netweaver profile, which doesn't reflect the optimizations needed for HANA → just check for yourself with "tuned-adm active", which profile is running
  - It offers a sap-hana profile, but this one contains also a lot of optimizations for x86 only ☹
  - For SLES 12 SP1+ the tool sapetune has to be installed and used to properly generate a SAP HANA on Power profile and to make it active –> SAP Note 2205917 & 1275776 describes the details

How to track SAP Note Changes?

- Use the favorites and notification function in the SAP ONE Support Launchpad

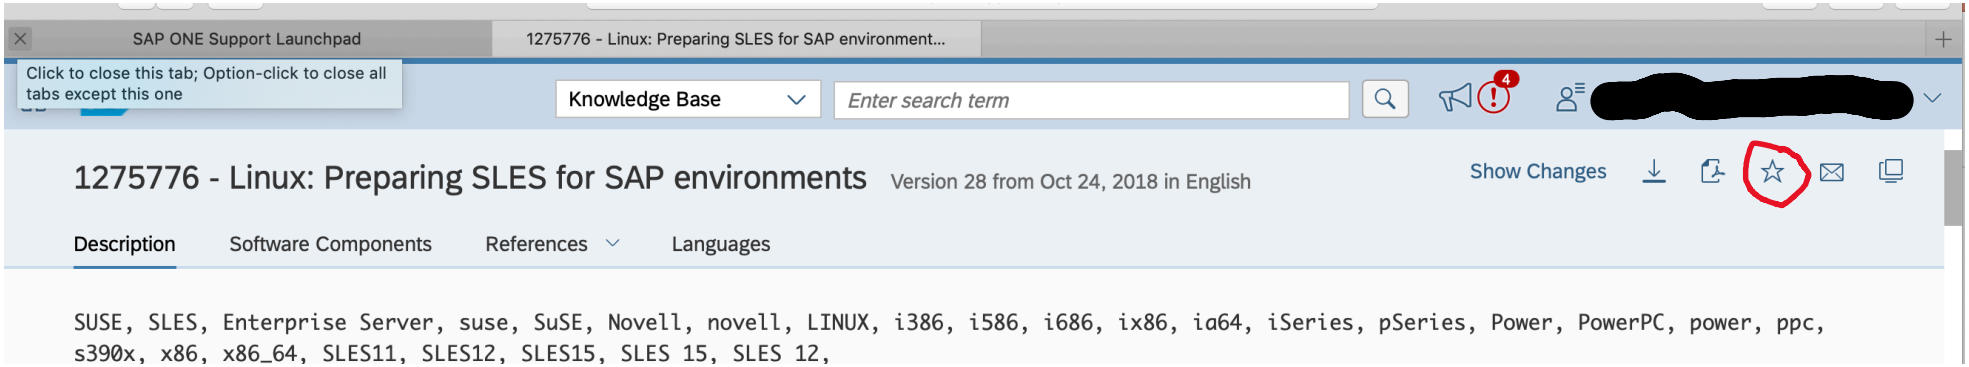

How to track SAP Note Changes?

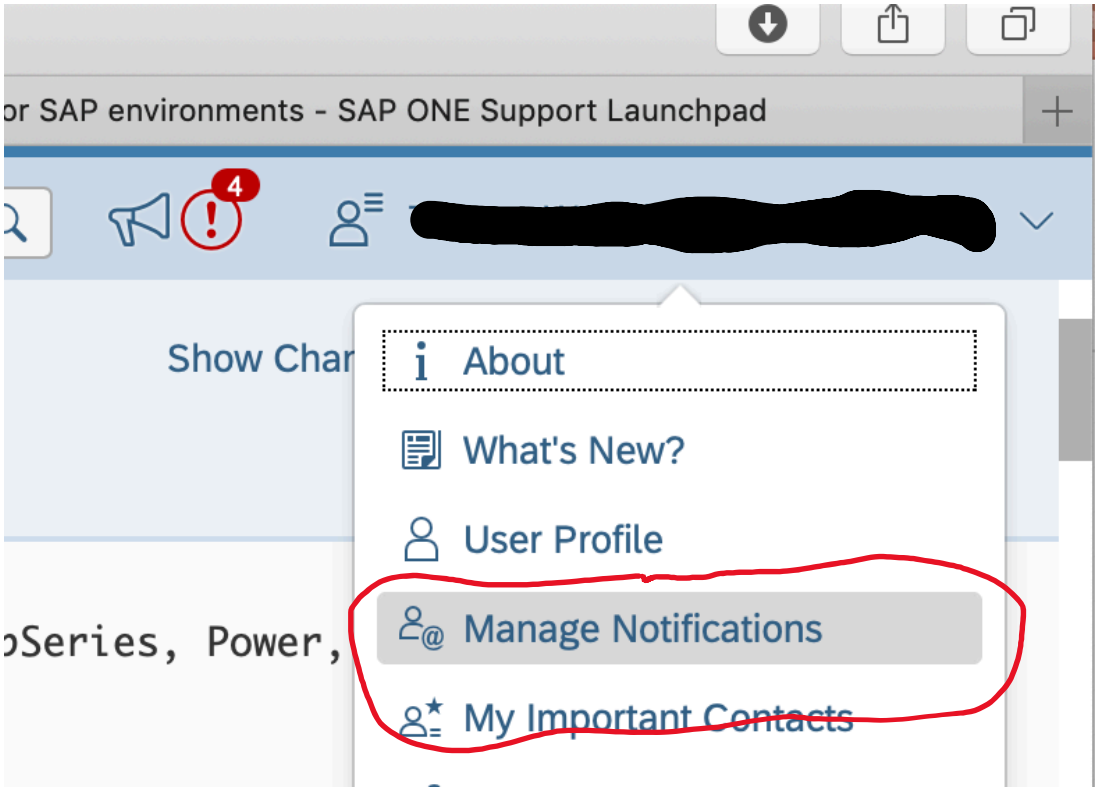

How to track SAP Note Changes?

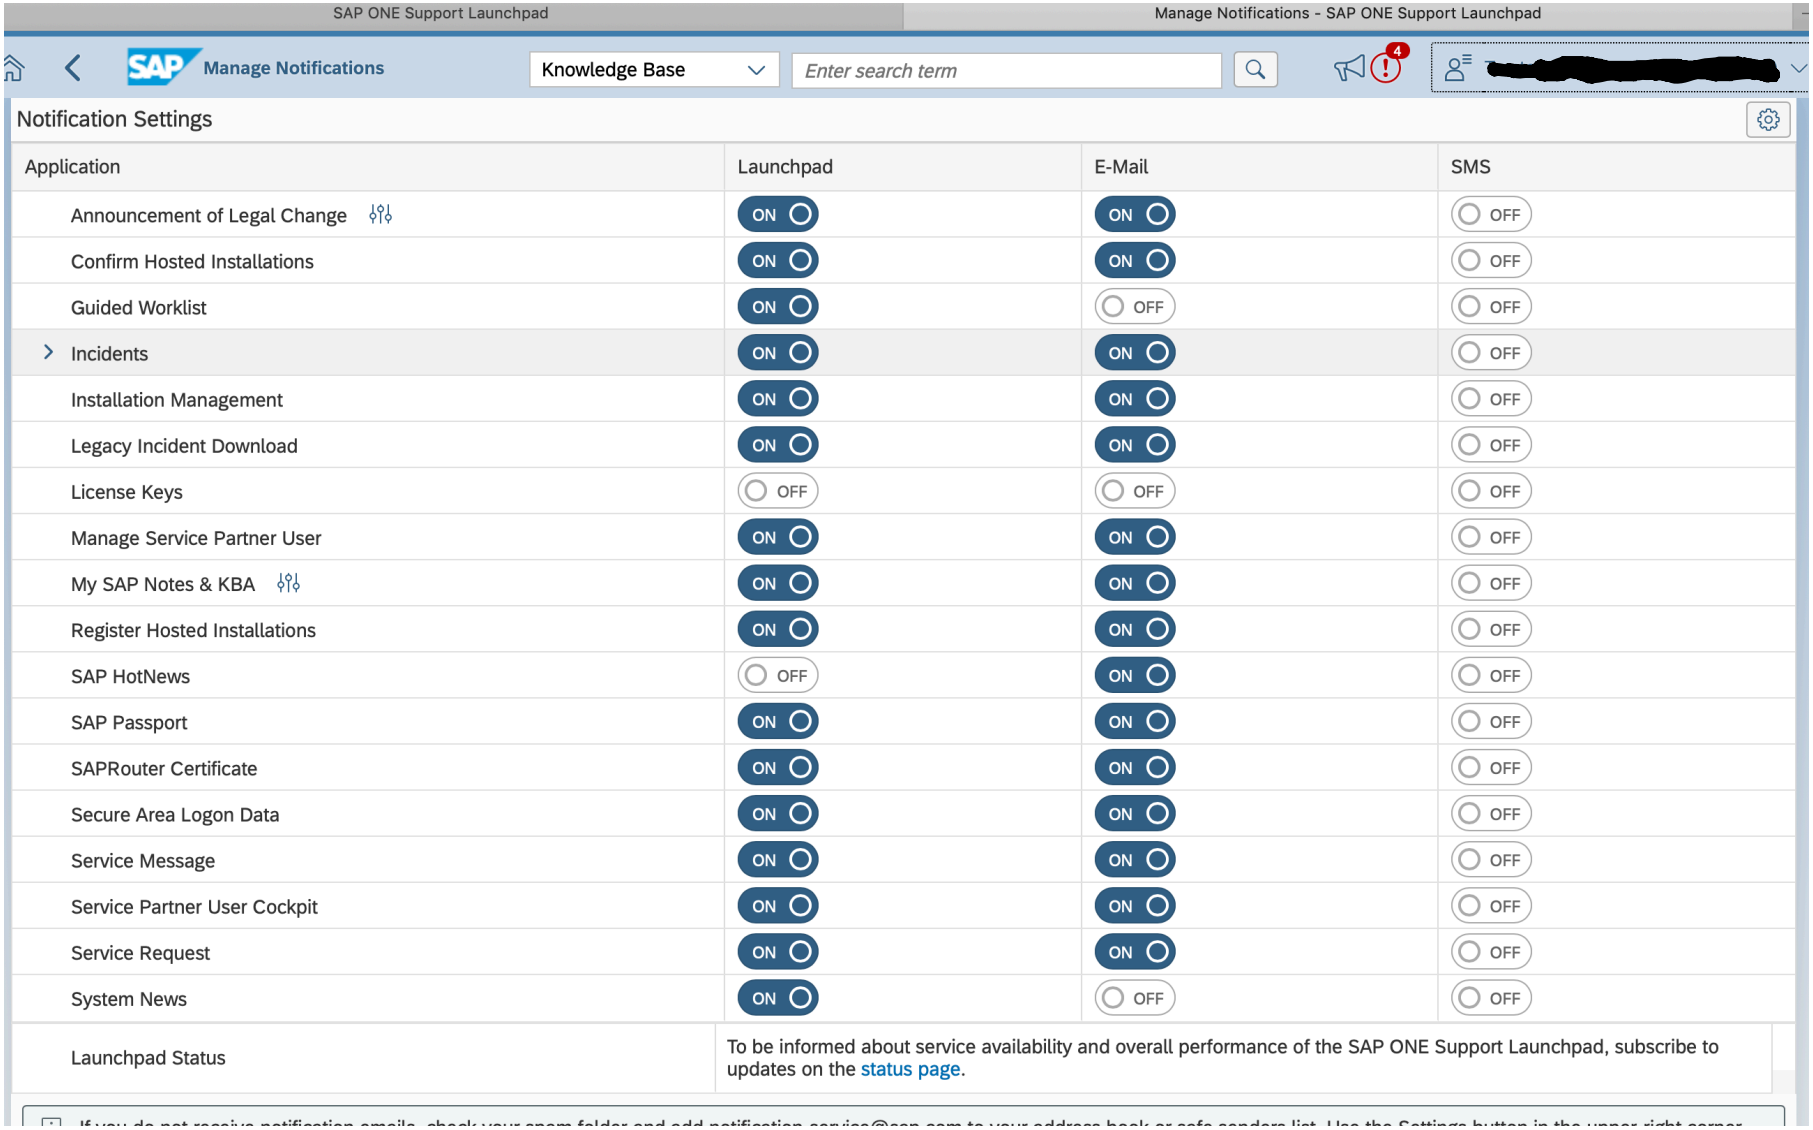

How to track SAP Note Changes?

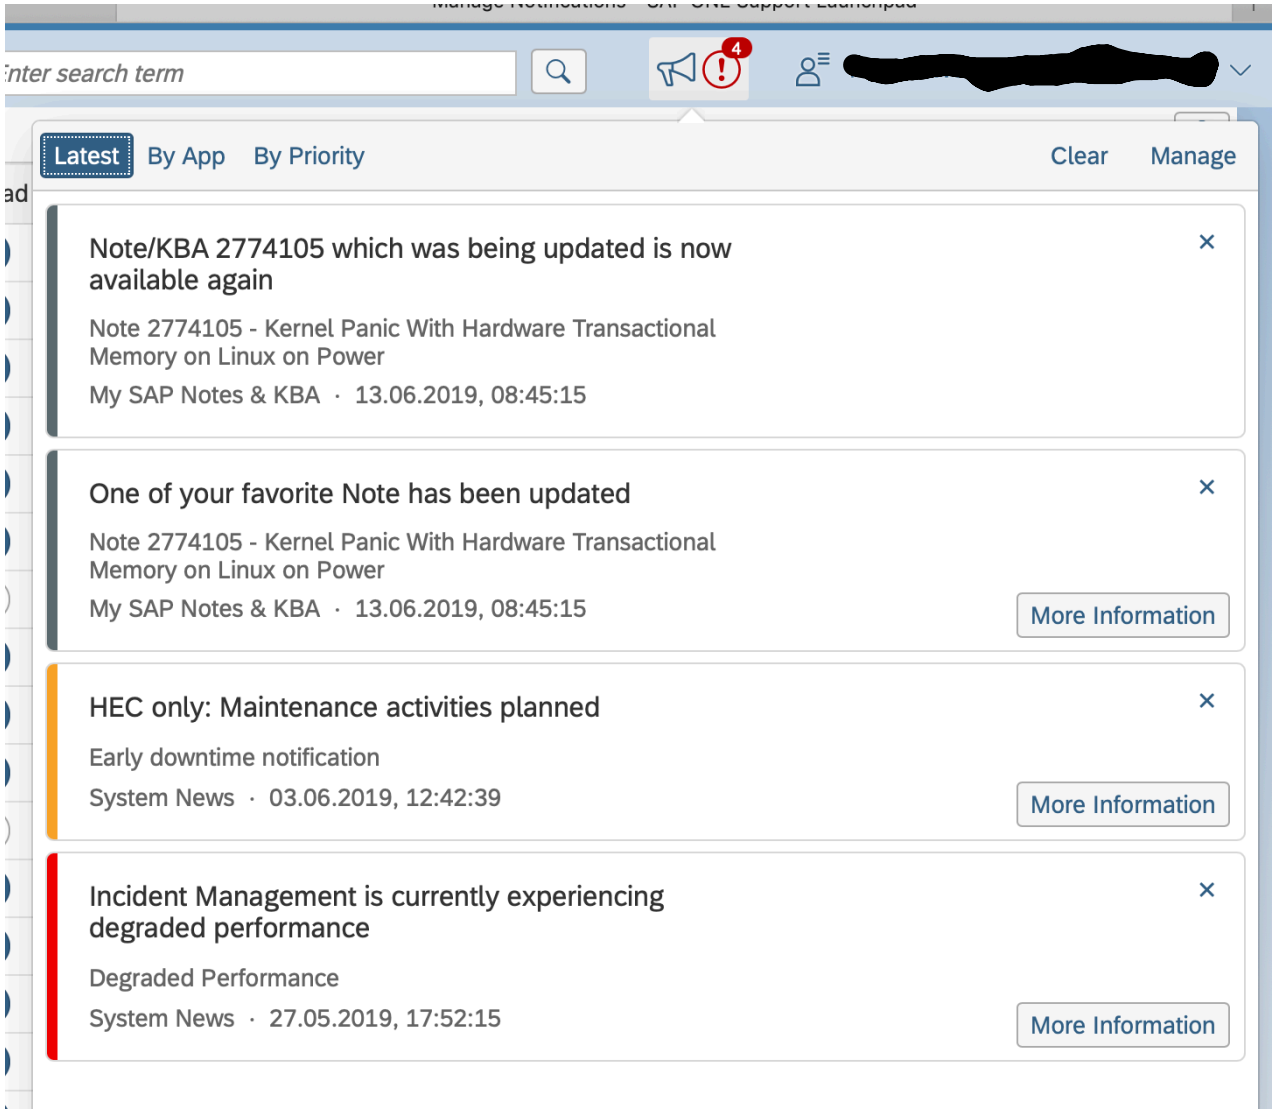

Optimize Your Code for HANA/OS Change

- Sometimes customers just migrating their DB to SAP HANA
  - Check your SAP Applications level about full support of HANA and migrate them too, if possible
  - Usually older SAP Applications levels supporting SAP HANA, but are not optimized → Sometimes you won't benefit much from HANA in terms of performance in such cases
- Self written code, specially batch jobs also need to be reviewed and optimized
- Specially customers which are switch from AIX to Linux (Power, or x86) with their ERP/ABAP have to review their job codes

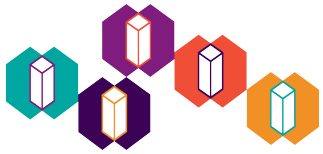

Optimize Your Shell Code for OS Change

- We ran some tests with a simple array script to simulate an customers ABAP job, which called a KSH script doing something similar like below

```
testmem: # Customers original version with “while”-loop
#!/bin/ksh93
l=0
lm=1000
while (( $l <= $lm ))
  do
    i=0
    im=4000
    while (( $i <= $im ))
       do
         array[$i]=abcdefghijklmnopqrstuvwxyzabcdefghijklmnopqrstuvwxyz$i
       (( i+=1 ))
       done
  (( l+=1 ))
done

testmem2: # Changed the code to use a ”for”-loop instead
#!/bin/ksh93
for l in {0..1000}
do
    for i in {0..4000}
    do
       array[$i]=abcdefghijklmnopqrstuvwxyzabcdefghijklmnopqrstuvwxyz$i
    done
done
```

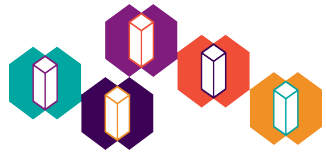

Optimize Your Shell Code for OS Change

```
testmem

AIX
# time ./testmem

  real    0m19.61s
  user    0m11.72s
  sys     0m0.00s

SLES 11 SP4
# time ./testmem

  real     0m27.719s
  user     0m27.710s
  sys 0m0.009s

SLES 12 SP1
  time ./testmem

  real     3m38.945s
  user     3m38.920s
  sys 0m0.000s
```

testmem2

- AIX:

```
# time ./testmem2

real    0m17.18s
user    0m9.18s
sys     0m0.95s
```

- SLES 11 SP4 (Power):

```
# time ./testmem2

real 0m17.829s
user 0m17.814s
sys  0m0.009s
```

The ABAP coding might suggest that you are testing CPU or memory speed, where in fact you are testing the efficiency of different areas of the ABAP interpreter and how well the platform compilers are able to optimize the ABAP processing loops. Keep in mind that SAP kernels are build with the XLC compiler on AIX, where as on Linux gcc is used.

Check, if your jobs using shell interpreter/script, specially KSH, before migrating from AIX to Linux for similar loop constructs and test them.

IBM SLS Power To Cloud Services for SAP HANA on Power

IBM SLS Power To Cloud Services for SAP HANA on Power

SAP HANA on Power Install

Overview

IBM experts will install SAP HANA on Power, which includes a pre-visit consultation on planning and design of the HANA on Linux architecture and layout.

Target Audience

- Clients deploying SAP HANA on Linux on Power
- All clients who are deploying a Production HANA landscape and require a certified installation of HANA on Power.

Benefits

- Results in a certified install of SAP HANA on Linux on Power

Qualifying Questions

- Are you planning to introduce SAP HANA to your environment?
- Have you chosen to run HANA on Power?

Key Features

- IBM Lab Services consultant will work with the client during pre-engagement planning to provide preliminary planning assistance with data center preparation, networking environment and solution environment prior to on site visit
- The on-site portion of the engagement may include the following activities:
  - –Create VIO Servers and setup NPIV and SEA adapters
  - –Installation of Linux
  - –Installation of SAP HANA Instance (Tailored Datacenter Integration)
  - –Verification of filesystem performance (HWCCT)
  - –Verification of Landscape
  - –Setup of SAP HANA replication

Deliverables

Post install documentation suitable for ongoing management and support of the system

Duration

The service varies depending on the size and complexity of the implementation, but can be customized to specific client requirements.

IBM SLS Power To Cloud Services for SAP HANA on Power

SAP HANA on Power Systems Health Check

Overview

IBM experts will assess the health of your HANA problems before they can impact your critical operations—including hardware, software and setup—to quickly identify areas of exposure and make recommendations.

Target Audience

- Clients who want to have the health of their HANA deployment checked
- Clients who want access to best practices and the latest technology to maximize their HANA investment and reduce risks

Benefits

- These services are designed to assess your SAP HANA software stack (virtualization, operating system, and SAP) and provide solution configuration and post-installation documentation.
- The result is a fine-tuned SAP HANA solution that can meet new and future demands.

Qualifying Questions

- Do you have a HANA installation?
- Has your HANA deployment changed and are you looking to verify best practices in your landscape?

Key Features

- Conduct a remote conference meeting between IBM and the client prior to the on-site visit to identify client's requirements
- Assess your SAP HANA solution
- Verify interoperability of the firmware levels, device drivers and OS patches
- Check service agent configuration
- Check for pending or open service events
- Verify OS and SAP HANA software based on SAP recommendations
- Update your hardware to the latest level of firmware, device drivers, OS patches, GPFS and SAP software as required on a mutually agreed-upon schedule
- Review VIO setup and look for potential bottlenecks
- Run SAP's HWCCT tool to verify filesystem performance and environment consistency
- Check error logs and status of hardware and software components
- Investigate and address any SAP HANA alerts
- Verify network connectivity and health
- Solicit client concerns and investigate as needed

Deliverables

Findings and recommendations

Duration

The service varies depending on the size and complexity of the implementation, but can be customized to specific client requirements.

IBM SLS Power To Cloud Services for SAP HANA on Power

SAP HANA on Power Performance Assessment

Overview

This offering is designed to assist clients running SAP HANA on IBM Power Systems get the best performance they can out of their HANA landscape. With this service, the system will be inspected by a consultant checking for both hardware and software configurations and system health conditions, with a focus on performance optimization.

Target Audience

- HANA on Power clients currently in production
- Clients looking to maximize performance for their HANA solution

Benefits

- Provide the client with recommendations to improve HANA system performance

Qualifying Questions

- Are you wanting to improve performance on your HANA installation?
- Do you want to validate best practices are used in your HANA installation?

Key Features

- Conduct an opening session with client's designated personnel to discuss:
  - Performance expectations
  - Current load
  - Areas of concern
- Evaluate Power system and software for SAP HANA where applicable
- Evaluate HMC configuration
- Evaluate partitioning–sizes, resources assigned
- Check firmware levels for currency
- Review VIO setup and look for potential bottlenecks
- Evaluate utilization of memory, disk, and paging space
- Evaluate CPU utilization/load
- Brief inspection of system logs for suspect error messages
- Verify network settings and network performance
- Use the SAP HWCCT tool to evaluate filesystem performance

Deliverables

Findings and Recommendations

Duration

Typical engagements will last one week but may extend longer depending on the size and complexity of the implementation.

IBM SLS Power To Cloud Services for SAP HANA on Power

SAP HANA on Power Migration Workshop

Overview

This service will provide a workshop environment to help the client understand the practical aspects of planning and executing a migration from any database or HANA on x/86 to HANA on Power. The workshop will discuss planning, sizing and executing the migration, SAP migration tools and methodologies, Best Practices and lessons learned from previous HANA migrations. A non-production sample migration can be added to the workshop and will provide hands on practical experience, guiding the client through each migration step.

Target Audience

- Clients running their current HANA solution on x/86 platforms.
- Clients running any database on x/86 and/or Unix systems wanting to move to HANA on POWER9.

Benefits

- Superior performance and reliability over x/86 systems
- Superior virtualization for flexibility and scaling
- Homogeneous POWER9 infrastructure for clients running other SAP workloads on POWER9

Qualifying Questions

- Do you currently run your HANA solution on x/86 platforms?
- Are you experiencing performance, reliability and scaling issues with your current HANA solution?
- Do you run other database environments like Oracle on x/86 and/or Unix environments and want to migrate them to HANA?

Key Features

- A review and assessment of the client's current SAP HANA environment to including database sizes and available outage windows.
- Migration tools and options covering homogeneous and heterogeneous environments.
- An assessment of potential risk areas.
- Business considerations….when can outages be schedule so they have minimal impact on operations.
- Programming considerations and need to update application code for UNICODE and HANATIZE transactions.
- Technical considerations including Backup/Restore and HA/DR and testing of these to certify production environment.
- Scheduling multi-landscape migrations…simultaneous or sequential.
- Downtime and planning the go-live weekend.
- Staffing…SAP Certified resources as well as testing and BASIS resources.
- Selection of a non-production environment for hands on work.
- Development of a high level plan that could be used for an eventual migration.

Deliverables

Workshop Summary Report
High level migration plan

Duration

The service duration varies depending on the size and complexity of the planned migration, but can be customized to specific client requirements such as to include a non-production sample migration.

IBM SLS Power To Cloud Services for SAP HANA on Power

SAP HANA on Power Advanced Features Deployment

Overview

This IBM Systems Lab Services offering is designed for clients planning to run SAP HANA workloads on IBM Power Systems. IBM Systems Lab Services provides on-site assistance in deploying additional HANA products and features, including High Availability and HANA Scale-out or cluster landscapes.

Target Audience

- Clients deploying SAP HANA on Linux on Power who wish to take advantage of advanced features in HANA such as HA and Scale-out

Benefits

- Delivered onsite by experienced IBM Lab Services consultants using proven methods and best practices
- Deployment with reduced risk of failures and issues due to incorrect or incomplete implementation
- Skills transfer from our experts helps you fully exploit the capabilities of this product

Qualifying Questions

- Are you planning to introduce SAP HANA to your environment?
- Have you chosen to run HANA on Power?
- Are you looking for a High Availability or Scale-out solution?

Team Contacts

Owner: Kurt Koehle – koehle@us.ibm.com

Key Features

- IBM Lab Services consultant will work with the customer during pre-engagement planning to provide preliminary planning assistance with data center preparation, networking environment and solution environment prior to on site visit.
- The on-site portion of the engagement may include one or more of the following activities
  - Setup of SAP HANA Replication
  - Setup of SLES SMT (Subscription Management Tool)
  - Setup of SLES-HA with HANA Replication
  - Setup of SAP HANA Host Auto Failover
  - Setup of SAP HANA Scale-out landscape
  - Setup of SAP HANA Dynamic Tiering
  - Setup of SAP HANA in a RedHat Pacemaker Cluster

Deliverables

Post install documentation suitable for ongoing management and support of the system

Duration

The service varies depending on the size and complexity of the implementation, but can be customized to specific client requirements.

IBM SLS Power To Cloud Services for SAP HANA on Power

SAP HANA on Power Linux Security Assessment

Overview

Database systems are by nature very popular targets for hackers and must therefore be protected. SAP HANA systems typically store business related information and are considered as being business critical.

The security of the underlying operating system is at least as important as the security of the SAP HANA database. Many hackers target the operating system in order to gain access to attack the running database application.

This assessment is a comprehensive analysis of your Linux security settings
for the specific purpose of protecting your SAP HANA database.

Target Audience

- Customers wanting to improve their Linux configurations for SAP HANA
- Customers wanting to stay abreast of the latest Linux security solutions
- Customers wanting a security baseline for defining standard builds for partitions running SAP HANA

Benefits

- Learn about security settings recommended by Linux for SAP HANA
- No sensitive data is collected in the security assessment
- The assessment only reads existing security settings – no settings are altered on the assessment partition

Qualifying Questions

- Do you have a HANA on Power installation with Linux?
- Do you want to stay abreast of the latest Linux security solutions?
- Are you looking for a security baseline for defining standard builds for partitions running SAP HANA?

Team Contacts

Owner: Stephen Dominguez – sdoming@us.ibm.com

Key Features

Phase 1 – Preparation (remote):
Conference calls are held prior to the service to validate the scope, agenda, schedule and required materials.
Example Tasks:

- Client provides overview of their current Linux Security environment
- IBM team prepares the service agenda/schedule
- IBM team details security data collection process
- IBM team provides customer security questionnaire
- Identify required materials / Finalize key players

Phase 2 – Review Findings (on-site):
Review the findings of the assessment with customer
Example Tasks:

- Consultant reviews the findings of the assessment with customer staff
- Customer reserves conference room with projector
- Customer staff can ask questions about the details of the findings
- Customer staff can ask questions about the security recommendations
- At the customer request, additional presentations can be provided to expound upon particular recommendations

Deliverables

- Security Assessment for SAP HANA Findings document
- Security Risk Heat Map
- Executive Summary (available upon request)

Duration

At least 1-3 days on-site (remote options are available)

Backup