

Managing Primo on a Large Scale

Stefan Lohrum, KOBV

IGELU conference Berlin 2013

Managing Primo on a Large Scale



- KOBV Primo Consortium
- Primo Infrastructure
- Disaster Recovery System
- Monitoring





ZIB & KOBV





- **Mathematics & Computer Science Research Fast Algorithms - Fast Computer**
- **II North-German Supercomputing Alliance**
- **Berlin Scientific Network**
- Data Storage & Archiving
- Regional Library Consortium (250+ Libraries)
- Public Services (Search Engine, ILL, Library Registry)
- Hosting Services (Primo, SFX, Metalib, Verde, OPUS, ...)



Primo Consortium – Audience



koby



Expectations

7 x 24 operation

Cost savings

Ease of operation

Learning from the consortial

partners





Working together



Open policy (``everybody can see everything´´), ...

... but partners work very carefully

BO tables become more and more confusing, ...

... but this may be a reminder from earlier versions

Pipe coordination required ...

... libraries seem to reload their catalogues every 3 months

Still deficits in multitenancy and robustness ...

... issues in one instance may have impact to others

Processes take long ...

... e.g. migration Primo 3 to Primo 4 took 7 month



=> No contact to user or patron, delayed error messages

=> Error fixing may require service provider or other 3rd party

=> Communication infrastructure required



Production system

2 FE, 4 SE, 1 BE/DB + 1 LB (each 8 cores / 32 GB RAM)

Staging system (t.b.d.)

2 FE, 4 SE, 1 BE/DB + 1 LB (12 cores / 144 GB RAM)

Disaster Recovery System

2 FE, 4 SE, 1 BE/DB + 1 LB (each 6 cores / 32 GB RAM)

Disaster Recovery System





=>> Business requirements

- =>> Cost-Benefit Analysis
- =>> Don't forget maintenance

Disaster Recovery System



ExLibris recommendations:

- (a) Have two separate Primo installations
 - Run all pipes on both systems
 - Use DP to copy user generated data from Prod to failover system
- (b) Have two separate Primo installations
 - Synchronize a set of directories (e.g. index file, SE conf. files)
 - Use Data Guard to synchronize the Oracle database

If you don't want to synchronize back the user generated data, have a CSS hiding e-Shelf and other functionality If failover system is off site, global load balancing is required



KOBV Disaster Recovery System



Model (a)

On demand clone database (backup / recovery)

Performed successful for 2 migrations

BUT:

Have to be very careful with configuration

Don't forget load testing !!!!

Primo Front End Infrastructure





Architecture, Components & Interfaces



- **PDS** & IDS (Identity Management System)
 - Third Node Retrieval (e.g. Primo Central),
- **Federated Search** (e.g. Metalib)
- **RTA** ("Run Time Availability")
- **OvP** ("OPAC via Primo")
 - Linking Service SFX
- **Front End Enrichments** (i.e. Google Books/Synectics covers)

=> Fault tolerance on outage of components?

=> How to find the error?

Stefan Lohrum, KOBV – Managing Primo on a Large Scale – IGELU – Berlin – 2013

System Weaknesses

In general:

Required Resources > **Given Resources** =>

Resource Limitations

Hardware: #Cores, RAM, IO, disk,

JBoss infrastructure # Threads

Keep in mind

- Load increases by usage
- **DOS attacks generate HUGE load**

=> Evaluating sizing is an ongoing process Baselines required









There are lots of ...

FE: library_server.log, localhost_access_log, lockedThread.log, maxThread.log, ...

SE: agent_9501.log, slice_1.log , ...

BE: publish_server.log, PrimoProcessExecutor.log, DeployAll.log, ...

Ever looked inside?

2013-09-13 07:31:16,394 ERROR [t-http-0.0.0.0-1701-269] [c-includeStaticHTML] - staticHTML () failed to load. (0 ms)org.apache.taglibs.

standard.tag.common.core.NullAttributeException: The "url" attribute illegally evaluated to "null" or "" in < import>

at org.apache.taglibs.standard.tag.common.core.ImportSupport.doStartTag(ImportSupport.java:133)

at org.apache.jsp.tiles.headerTile_jsp._jspx_meth_c_005fimport_005f0(headerTile_jsp.java:5765)

at org.apache.jsp.tiles.headerTile_jsp.jspx_meth_c_005fcatch_005f1(headerTile_jsp.java:5730)

at org.apache.jsp.tiles.headerTile_jsp._jspService(headerTile_jsp.java:709)

at org.apache.jasper.runtime.HttpJspBase.service(HttpJspBase.java:70)

at javax.servlet.http.HttpServlet.service(HttpServlet.java:717)

at org.apache.jasper.servlet.JspServletWrapper.service(JspServletWrapper.java:369)

at org.apache.jasper.servlet.JspServlet.serviceJspFile(JspServlet.java:322)

=> There are many, mostly undocumented

=> Very unclean, mugged with stack traces

=> Essential information missing, e.g. IP if behind a load balancer

Monitoring

kobv

Audience

not only IT or Primo management

Alerts to all devices

Workstations, PCs, Tablets, Smartphones, ...

Monitoring of the whole infrastukture

Primo is NOT only Primo software

Active Information in case of problems

Escalation strategy





Monitoring

kobv

Error Tracking Support

Support of an iterative process Learning from errors...

Proactive Information ... and error detection ahead of problems

"End to End" Monitoring … the user's view on the system

Documentation of the availability

=>> Monitoring is part of a process







Most important parameters

- Pipes
- Processes (Indexing, Hotswapping, DYM)
- Search Performance
- Tablespaces
 Disk Free
 CPU Load







- Most important parameters
- System components

- Monitor Primo Status
- Pipe Monitoring Monitor the pipe processes and view errors
- Tools Monitoring Monitor the tools processes and view errors
- Process Monitoring Monitor the processes and view errors
- Dob Monitoring Monitor all the jobs running on the system
- Deploy Monitoring Monitor Deploy Jobs
- Search Engine Monitoring Monitor and manage search servers and indexing
- Watchdog Monitoring Monitor and manage Primo Watchdog

kobv

- Most important parameters
- System components

Process Name	Process Status	Execution Date	Execution Time	Records Succeeded	Records Failed	More
Hotswapping	Completed	09.01.2013 13:09:04	00:41:20	0	0	History Clean-up Execute
Indexing	Completed	09.01.2013 09:45:56	00:56:44	177001	0	History Clean-up Execute
Indexing_and_Hotswapping	Stopped	20.02.2013 23:00:00	00:02:11	7000	0	History Clean-up Execute
Didumean	Completed	10.06.2012 12:55:05	00:33:14	0	0	History Clean-up Execute
Indexing_and_Didumean_and_Hotswapping	Completed	16.02.2013 23:00:00	02:07:01	97558	0	History Clean-up Execute
Process Status : All						
Go back						Refresh

Tabelle: M_P_PROCESS_CONTEXT

kobv

- Most important parameters
- System components
- But somtimes
 - plugins required





NAGIOS Excurse

- Full-blown, server based monitoring system
- Sensors for all system parameters
- Public domain, package
- Schedules
- Service hierarchies, views
- Escalation schemas
- Protocol, trends etc.
- Browser interface, Android and IOS Apps
- Plugin interface



kobv

NAGIOS Excursus - Plugin Interface



- SSH Acces to external target systems (using a certificate)
- Shell script evaluates performance data ("sensor")
- Protocolling on Nagios
 Server (MySQL data base)
- Oracle plugin can run
 SQL commands
- Screen scraping using perl library "mechanize"

```
#!/sbin/sh
# check if aleph-union-process is running
                            | if online
# Aleph20
            0 | OK
                            | if disabled
            1 | WARNING
            1 | WARNING
                            | if maintenance
                            | if offline
            2 | CRITICAL
            3 | UNKNOWN
                            | if not available
if [ $# -lt 2 ]
        echo "usage: `basename $0` <host> <user>"
then
        exit 3
fi
HOSTNAME=$1
USR=$2
USR CMD="ps -f -u $USR | egrep 'b union|p union'"
CMD=/usr/bin/ssh
TEMPFILE=/usr/local/icinga/var/temp sshproc.$$
( $CMD icinga@$HOSTNAME $USR CMD > $TEMPFILE ) &
pid=$!
wait $pid
cnt=`cat $TEMPFILE | wc -l`
if [ $cnt -qt 0 ]
then
        stat=0
        echo "$stat - OK: (b union p union) x $cnt instances are
               running"
else
        stat=1
        echo "$stat - Warning: (b union p union) is NOT running"
fi
rm -f $TEMPFILE
exit $stat
```



Standard sensors for CPU load, IO wait, uptime, disk usage, free memory

- Process status, search status
- Logfile monitoring for typical error messages
 - (e.g. "PDS not reachable", "Timeout after 30000ms",…)
 - PDS response times

Penetration ("End to End") of Front End,

internal and external (Pingdom)

Additional sensors (e.g. garbage collection) as necessary

Primo Monitoring with Nagios



			Host 🔭		Service **		Status 🔭	Last Check 🛶		
			vs30.kobv.de	P	CPU	<u></u>	ОК	10-09-2013 01:43:49		
					CPULOAD	ૢૼૢૺ	ОК	10-09-2013 01:42:10		
					DISK	<u></u>	ОК	10-09-2013 01:44:38		
	Primo (h	g0010)			МЕМ	્ંં	ок	10-09-2013 01:43:27		
Host	Status	Services	Actions		PING	<u></u>	ок	10-09-2013 01:43:11		
kobv1.tubit.tu-berlin.de	UP	1 OK			PRIMO-vs30-CheckSystemBusyLog	्र ध अ ह	ОК	10-09-2013 01:45:13		
vs30.kobv.de	UP	12 OK			PRIMO-vs30-PDS-FUB-URL	<u></u>	ок	10-09-2013 01:44:38		
ve32 koby de	IID	5.0K			PRIMO-vs30-PDS-HUB-URL	ૢૺ૽	ОК	10-09-2013 01:44:59		
V352.K00V.06	UF	JOK			PRIMO-vs30-PDS-MAN-URL	್ಷಣ	ОК	10-09-2013 01:43:57		
vs33.kobv.de	UP	5 OK			PRIMO-vs30-PDS-TUB-URL	ું	ок	10-09-2013 01:44:47		
vs34.kobv.de	UP	5 OK			PRIMO-vs30-PDS-ULD-URL	<u></u>	ок	10-09-2013 01:45:28		
vs35.kobv.de	UP	5 OK	📄 🚞 🗋 📩		PRIMO_KOBV-vs30-pingdom-URL	ૢૺ	ОК	10-09-2013 01:41:03		
vs36.kobv.de	UP	1 OK						12 Matching		
vs37.kobv.de	UP	11 OK			vs30.kobv.de / PRIMO-vs30-PDS-ULD-URL					
vs39.kobv.de	UP	5 OK			3.0 2.0 1.0	mlilearre	have	hera Mareda_		
					Mon]rtime 817.9722 m Last	12:00 3366.9722	m Max 927.08	Tue 00:00 99 m Average Default Template		

End to End Monitoring



User experience with Primo

Search local repository (e.g. "berlin") Check no. of hits (should be > 10)

Error on timeout or not enough

External provider (pingdom.com)

Integrated statistics







Summary



Primo multitenancy is still limited, robustness

- Primo monitoring using Backoffice is not sufficient
 - NAGIOS and specific plugins for Primo allow monitoring, also in a hosted environment
- Monitoring is part of an "integrated" process

@ExLibris:

Have APIs for monitoring

Cleanup logiles, document them, consider load balancer

Questions?





Contact: Stefan Lohrum lohrum@zib.de