SC09 (http://www.exascale.org). IESP ambitiously intends to organize international agencies around sharing costs and leveraging each other's work in the short, medium, and long term. Tony Drummond gave a review of the DOE ACTS Collection, which curates and educates in the area of HPC software.

The Function (and Malfunction) of Software in a Production Computing Environment

In order to set the scope of this workshop, which was mainly attended and sponsored by those involved in procuring and managing high performance computing centers, we first had to define the type of software we are addressing. This workshop divided the HPC software ecosystem into the following strata:

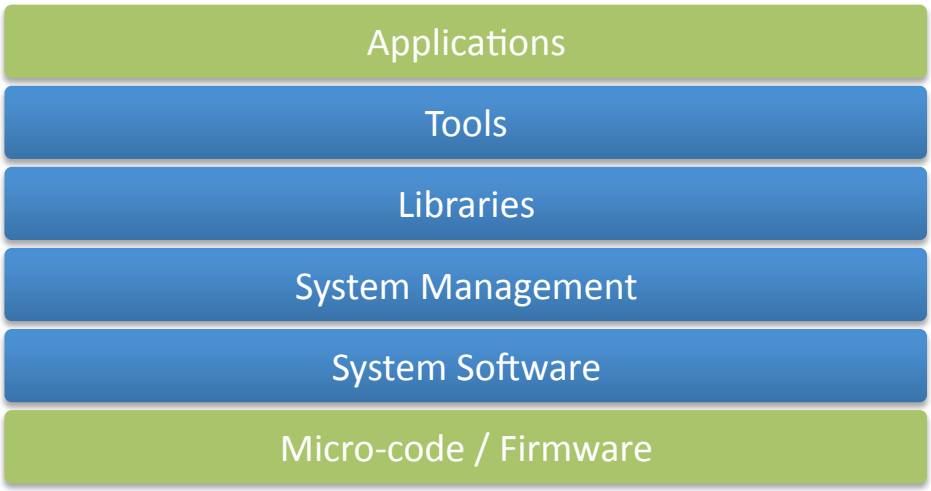

and identified the middle four layers as the scope of what HPC managers might best address as software best practices. The top and bottom software functions are, at many HPC centers, beyond the scope of managers or staff to control. Applications often come already developed from our customers. Very low-level software is often outside the control of those who purchase, rather than develop, HPC hardware. These circumstances vary between centers, and our choice to limit the scope of what we mean by *software* is meant to balance the discussion on common ground rather than to exclude topics outside the four areas above. Certainly the topic of HPC best practices in application-level software could comprise in itself an equally large or larger set of discussions.

This software stack's top-level function, through the execution of application code, is to make computing platforms a useful tool for science. The underlying layers, which should function reliably with good performance, all serve that end goal, and the degree to which they function depends largely on integration and maintence of software. An a priori postion of the attendees, based on decades of experience within HPC, is that both hardware and software systems often do not function out-of-the-box (http://www.nersc.gov/projects/HPC-Integration/) but instead require significant attention at all stages of the software lifecycle.

Software lifecyles vary from faciltiy to facility. This workshop categorized four phases as distinct parts in the HPC software lifecycle.

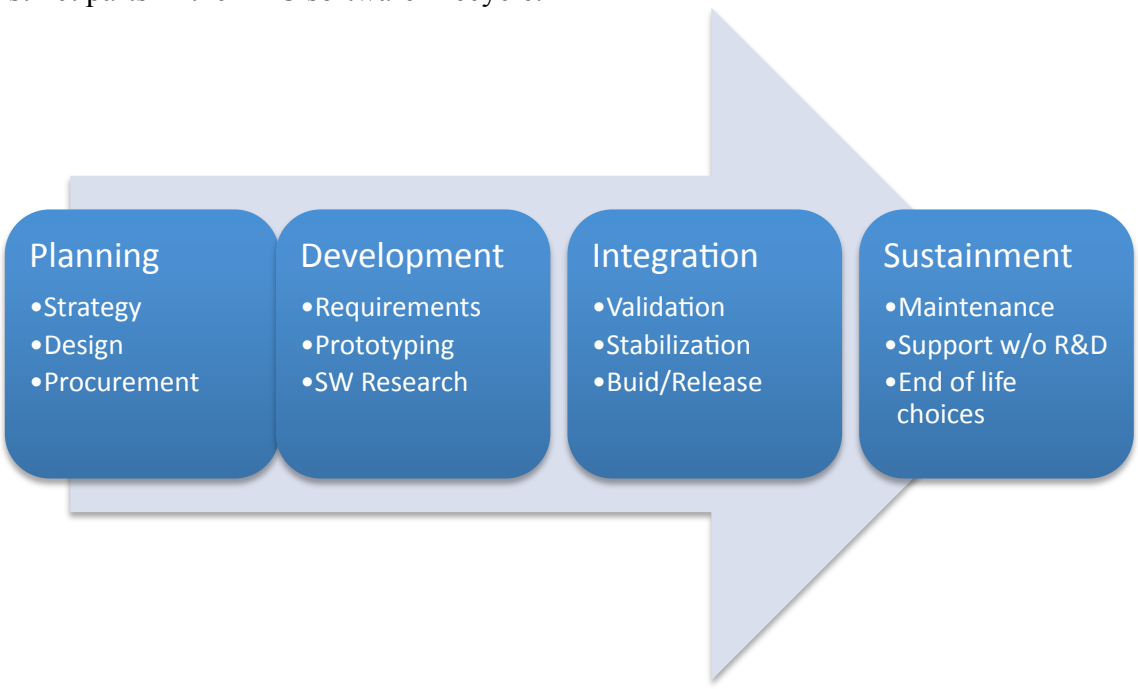

The term *production* is often used in a rather loose way to connote a stable computing platform which experiences infrequent failures or interruptions. Managers of HPC centers most often take a metric-based approach to indentifying the frequency, duration, and cause of interruptions in service. Much of the discussion and body of this report focuses on practices which drive down malfunction rates in HPC software.

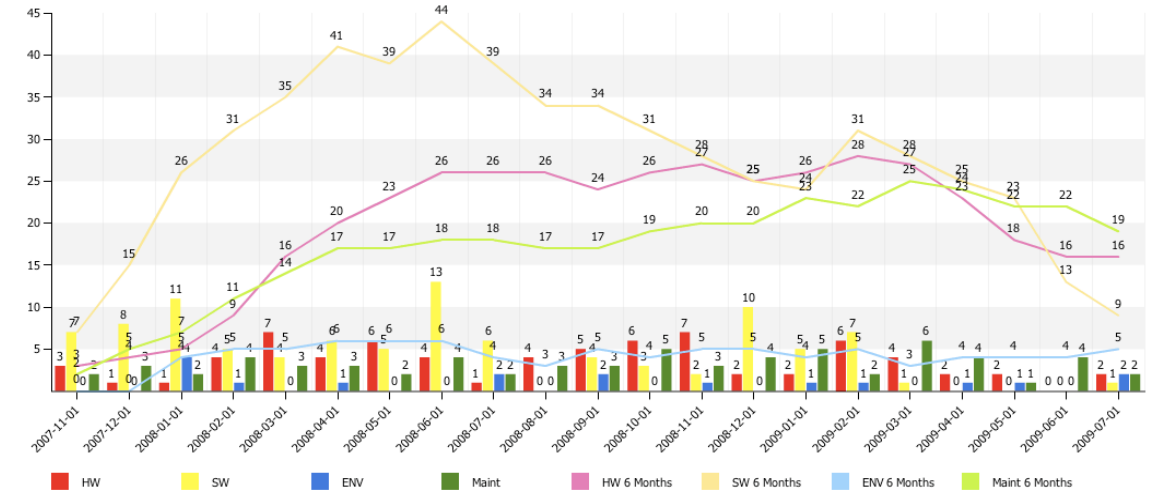

System outages for a 20K core machine over two years sorted by outage type.

In short, when software does not function, HPC does not function. Increasingly the preponderance of interruptions in HPC systems are caused by software faults rather than

Sys Software

MPI, MPT Based on MPICH2 version 1.0.2	1		1	1	1		1	3		1	1
PowerMan, FreeIPMI Power Management					1		2	1		3	-
genders, dist Configuration Management Tools					-		3			3	-
lm_sensors, FreeIPMI HW Sensors Support					1		3	3		3	-
lxbios, cmos_util CMOS/BIOS utilities					3		3			3	3
whatsup, skummee, cerebro Host Monitoring					-		3				1
yaci, gedi cluster installation					1		3				3
ANL UserBase					S						-
ARMS ANL RAS Management Suite					S						-
Asset Tracker Breakfix Asset Tracking					-					1	-
CHAOS See below for RAS components					-		1			3	-
CSA											-
Cacti Monitoring / Security					S			2			3
Catamount					-			1		1	-
Cbench Scalable Benchmarking and Testing Framework					-					1	-
Cobalt-Accounting					S						1
Cobalt-ResMangr					S						1
Cobalt-Scheduler					S						1
ConMan Console Management					S		2	1		3	3
DIM						1					-
EDAC Hardware Error Detection/Correction							1			3	-
GPFS system wide filesystem			1		1	1	3	1			1

Sys Software

	NICS	LANL	Blue Waters	DoD Create	ANL	BSC-CNS	LLNL	NERSC	ORNL	SNL	Unknown
Ganglia Adapted to MareNostrum in order to improve scalability and include other fetaures not supported			1		-	1		1			1
IBM CNK					1		1				1
IBM ION					S		1				-
LDAP System wide auth and unix groups	1	1	1	3	S		1	1			3
LMT Lustre Monitoring Tool					-		3	1		3	3
Linux	1			1	1		1	1		1	1
Lustre Forwards requests to Lustre servers on service nodes					-		1	1		1	1
MOAB Tri-lab Workload Manager	1	1			-		1	1			-
MPI "IB, MVAPICH, MVAPICH2 based from Ohio State"		2	1	2	-		1	1		1	1
MPI-IO Part of MPICH2-based MPI from Cray	1		1		1		2	1		2	-
MX					1	1					-
Moab Tri-lab Workload Manager		1			-	1	1	1	2	2	3
Munge Scaleable Cluster Authentication					-		3			3	3
NFS "Support included in Linux 2.6 kernel distribution; NFS v3 support, not sure about NFSv4"		1	1	1	1	1	1	1		1	1
Nagios Monitoring / Security	1				S	1	3	2			1
OFED/OpenIB		1			1		1	1		1	1

OpenSsh (with OTP extensions)	system access					1		1	1		1	1
PBS Pro					3	-			1	2	2	2
PVFS	Parallel file system					1			1			2
Pam	Used widely for authentication					1		1	1		1	1
Portals	Based on Portals from Sandia					-			1		1	-
PowerMan	Power Management					1		3				3
RASilience	Breakfix Asset Tracking										1	-
RT			3			S		1			2	-
Red Hat Linux	CHAOS built on top of Red Hat								1		1	1
Request Tracker	Breakfix Asset Tracking					-					1	-
SHMEM	Part of xt-mpt			1		-			1		2	2
SLES					3	1	1		1		1	1
SLURM	Highly Scalable Resource Manager		3			-	1	1			1	R
SUSE	on BG/L login and service nodes					1		3				1
Sun HPC Stack	Tools and configurations for Sun cluster management					-					1	R
Torque		1	1			-			1		2	R
Trac						3			1		1	-
ZeptoOS-CN						R			2			R
ZeptoOS-ION						R, S						R
cfengine	Configuration Management Tools		3	1		-		3	1			-
crms	"Have a module file, but references non-existent directory"									2	2	-
gPXE	Boot-over-IB Support for PXE Boot				3						1	1
ldap			1	1		1		1			1	1
mrsh	Scalable rsh implementation					-		3			3	-
oneSIS	Diskless image management system					-					1	-
pam				1		S		1			1	3
pdsh	Parallel Shell					S		3	3		1	-
ssh				1		1		1	1			1
yaci	cluster installation		1		2	-		3				-
jira					1							
Romio						1						
PXE						S						
Zenoss			1									
OpenMPI			1									
puppet		1										