



CHAPTER 3

Performance and Health Monitoring

This chapter is organized as follows:

- [Collected Performance Data, page 3-1](#)
- [Monitored Health Metrics, page 3-10](#)

Collected Performance Data

This section summarizes the performance data collected by the collectd monitoring agent which is installed on all nodes. While some of the collected system-specific performance data is common for all nodes (for example disk space, CPU), the collectd agent uses plug-ins to collect application-specific data (for example for MBean, Tomcat, Apache).

This data can be accessed in several ways:

- From the Director UI > System > Stats.
- Through the WebEx Social API.

Table 3-1 *Collected Performance Data*

Type	Instance	Matrix	Description	Role
CPU	core#	idle	Percentage of time that the CPU or CPUs were idle and the system did not have an outstanding disk I/O request.	All
		interrupt	Percentage of time spent by the CPU or CPUs to service hardware interrupts.	
		nice	Percentage of CPU utilization that occurred while executing at the user level with nice priority.	
		softirq	Percentage of time spent by the CPU or CPUs to service software interrupts.	
		steal	Percentage of time spent in involuntary wait by the virtual CPU or CPUs while the hypervisor was servicing another virtual processor.	
		system	Percentage of CPU utilization that occurred while executing at the system level (kernel). Note that this does not include time spent servicing hardware and software interrupts.	
		user	Percentage of CPU utilization that occurred while executing at the user level (application).	
		wait	Percentage of time that the CPU or CPUs were idle during which the system had an outstanding disk I/O request.	

Table 3-1 Collected Performance Data (continued)

Type	Instance	Matrix	Description	Role
Disk Usage	boot	used	Used space on partition /boot	All
		reserved	Space on /boot partition reserved for root user.	
		free	Free space on partition /boot	
	opt	used	Used space on partition /opt	All
		reserved	Space on /opt partition reserved for root user.	
		free	Free Space on /opt partition.	
	root	used	Used space on partition /	
		reserved	Space on /opt partition reserved for root user.	
		free	Free Space on /opt partition.	
Disk	sda/sda1/sda2/sdb	disk_merged read	The number of read operations, that could be merged into other, already queued operations, i. e. one physical disk access served two or more logical operations.	All
		disk_merged write	The number of write operations, that could be merged into other, already queued operations, i. e. one physical disk access served two or more logical operations.	
		disk_octets read	Bytes read from disk per second	
		disk_octets write	Bytes written to disk per second	
		disk_ops read	Read operation from disk per seconds	
		disk_ops write	Write operation to disk per seconds.	
		disk_time read	Average time an I/O- read operation took to complete, equivalent to svctime of vmstat	
		disk_time write	Average time an I/O-write operation took to complete, equivalent to svctime of vmstat	
Disk Usage	boot, opt, root	free	Used space on a specified partition.	All
		reserved	Space on a /opt partition reserved for root user.	
		used	Free space on a specified partition.	
DNS	octets	queries	Number of octets sent.	All
		responses	Number of octets recieved	
	opcode	opcode9	Number of packets with a specific opcode, e. g. the number of packets that contained a query.	All
		query	TBD	
	qtype	#0	Number of queries for each record type #0.	All
		a	Number of queries for each record type a.	
		aaaa	Number of queries for each record type aaa.	
		ptr	Number of queries for each record type ptr.	
		txt	Number of queries for each record type txt.	
Interface	eth0	if_errors rx	Rate of Error in receiving data by network interface.	All
		if_errors tx	Rate of Error in transmitting data by network interface.	
		if_octets rx	Rate of Bytes received by network interface.	
		if_octets tx	Rate of Bytes transferred by network interface.	
		if_packets rx	Rate of packets receivedby network interface	
		if_packets tx	Rate of packets transferred by network interface	
	lo	if_errors rx		All
		if_errors tx		
		if_packets tx		

Table 3-1 Collected Performance Data (continued)

Type	Instance	Matrix	Description	Role
Load		longterm	longterm represents the average system load over 15 min period of time.	All
		midterm	midterm represents the average system load over 5 min period of time.	
		shortterm	shortterm represents the average system load over 1 min period of time. Refer top/w/uptime man page for more details.	
Memory		buffered	The amount of memory used as buffers.	All
		cached	The amount of memory used for caching.	
		free	The amount of idle memory.	
		used	The amount of memory used Refer free/vmstat man page for more details.	
NTP	frequency_offset	loop		All
	time_dispersion	local		All
		<NTPServer>	Value indicates the magnitude of jitter between several time queries in MS	
	time_offset	error		All
		loop		
		<NTPServer>	Value shows the difference between the reference time and the system clock in MS	
	delay	<NTPServer>	Value is derived from the roundtrip time of the queries in MS	All
Swap	swap	cached	Memory that once was swapped out is swapped back in but still also is in the swapfile (if memory is needed it doesn't need to be swapped out AGAIN because it is already in the swapfile. This saves I/O) (http://www.redhat.com/advice/tips/meminfo.html/)	All
		free	Total amount of swap space available.	
		used	Total amount of swap space used	
	swap_io	in	Amount of memory swapped in from disk	All
		out	Amount of memory swapped out from disk	
Uptime		uptime	Second since VM is running.	All

Table 3-1 *Collected Performance Data (continued)*

Type	Instance	Matrix	Description	Role
VMWare	CPU	elapsed_ms	Retrieves the number of milliseconds that have passed in the virtual machine since it last started running on the server. The count of elapsed time restarts each time the virtual machine is powered on, resumed, or migrated using VMotion.	All
		limit_mhz	Retrieves the upper limit of processor use in MHz available to the virtual machine.	
		reservation_mhz	Retrieves the minimum processing power in MHz reserved for the virtual machine.	
		shares	Retrieves the number of CPU shares allocated to the virtual machine.	
		stolen_ms	Retrieves the number of milliseconds that the virtual machine was in a ready state (able to transition to a run state), but was not scheduled to run	
		used_ms	Retrieves the number of milliseconds during which the virtual machine has used the CPU. This value includes the time used by the guest operating system and the time used by virtualization code for tasks for this virtual machine. Percentage of cpu utilization is $\text{used_ms} * \text{number_of_core} / \text{elapsed_ms}$	
	Memory	active_mb	Retrieves the amount of memory the virtual machine is actively using—its estimated working set size	All
		ballooned_mb	Retrieves the amount of memory that has been reclaimed from this virtual machine by the vSphere memory balloon driver (also referred to as the vmmemctl driver)	
		limit_mb	Retrieves the upper limit of memory that is available to the virtual machine.	
		mapped_mb	Retrieves the amount of memory that is allocated to the virtual machine. Memory that is ballooned, swapped, or has never been accessed is excluded	
		reservation_mb	Retrieves the minimum amount of memory that is reserved for the virtual machine	
		shares	Retrieves the amount of physical memory associated with this virtual machine that is copy-on-write (COW) shared on the host.	
		swapped_mb	Retrieves the amount of memory that has been reclaimed from this virtual machine by transparently swapping guest memory to disk	
		used_mb	Retrieves the estimated amount of physical host memory currently consumed for this virtual machine's physical memory	
Apache		apache_connections		App Server & Worker
		apache_idle_workers		
	apache_scoreboard	closing		App Server & Worker
		dnslookup		
		finishing		
		idle_cleanup		
		keepalive		
		logging		
		open		
		reading		
		sending		
		starting		
		waiting		

Table 3-1 *Collected Performance Data (continued)*

Type	Instance	Matrix	Description	Role
State Manager	StateManager HTTP Response Code	activemq-code		App Server & Worker
		cache-code		
		digest-code		
		graph-code		
		index-code		
		json-code		
		notifier-code		
		quad-code		
		quad_analytics-code		
		rabbitmq-code		
		rdbms-code		
		recommendation-code		
		search-code		
Processes	fork	fork_rate	Number of new process forked per second.	All
	ps_state	blocked	Count of processes in Blocked state. If consistently high, alert condition need attention.	All
		paging	Count of processes in Paging state. If consistently high or growing, alert condition need attention.	
		running	Count of processes in running state. Typically less or equal to num of cores.	
		sleeping	Count of processes in sleeping state. Typically most processes are in this state.	
		stopped	Count of processes in Stopped state	
		zombies	Count of processes in Zombies state. If consistently high or growing, alert condition need attention.	
TCP Connection	Port 80 - App Server, Port 80 - Worker, Port 80 - Director-Web, Port 61616 - Message Queue, Port 8983 - Search Store, Port 7973 - Index Store, Port 27001 - Analytics Store, Port 27000 - JSON Store, Port 11211 - Cache	close_wait		App Server, Worker, Director-Web, Message Queue, Search Store, Index Store, Analytics Store, JSON Store, Cache
		closed		
		closing		
		established		
		fin_wait1		
		fin_wait2		
		last_ack		
		listen		
		syn_recv		
		syn_sent		
		time_wait		

Table 3-1 Collected Performance Data (continued)

Type	Instance	Matrix	Description	Role
Oracle		blockingLock		RDBMS Store, Graph Store
		cacheHitRatio		
		dbBlockBufferCacheHitRatio		
		dictionaryCacheHitRatio		
		diskSortRatio		
		invalidObjects		
		latchHitRatio		
		libraryCacheHitRatio		
		lock		
		lockedUserCount		
		offlineDataFiles		
		pgaInMemorySortRatio		
		rollBlockContentionRatio		
		rollHeaderContentionRatio		
		rollHitRatio		
		rollbackSegmentWait		
		sessionPGAMemory		
		sessionUGAMemory		
		sgaDataBufferHistogramRatio		
		sgaSharedPoolFree		
		sgaSharedPoolReloadRatio		
		softParseRatio		
		staleStatistics		
	ioPerTableSpace: ecp_data, sysaux, system, undotbs1, users	PHY_BLK_R		RDBMS Store, Graph Store
		Phy_BLK_W		
	oraUsageTablespace: ecp_data, sysaux, system, undotbs1, users	free_mb		RDBMS Store, Graph Store
		percent_free		
		percent_used		
		size_mb		

Table 3-1 Collected Performance Data (continued)

Type	Instance	Matrix	Description	Role
Solr	Search	avgRequestsPerSecond	Number of requests server per second	Search Store
		avgTimePerRequest	average time taken to server each request	
		errors	Rate of error, requests that returned error.	
		requests	Rate of request servered by SOLR.	
		timeouts	Rate of request timed out, request that failed due to time out error.	
	Search: documentcache, fieldvaluecache, filtercache, queryresultcache Index: autocompletefieldvalue, followerfieldvaluecache, postfieldvaluecache, socialfieldvaluecache, videofieldvaluecache	cumulative_evictions		Search Store, Index Store
		cumulative_hits		
		cumulative_inserts		
		cumulative_lookups		
		evictions		
		hitratio		
		hits		
		inserts		
		lookups		
		size		
		warmupTime		
	Search: searcher Index: autocomplete, follower, post, social, video	maxDoc		Search Store, Index Store
		numDocs		
Java Memory		HeapMemoryUsage_committed		Search Store, Index Store, Message Queue, App Server, Worker
		HeapMemoryUsage_init		
		HeapMemoryUsage_max		
		HeapMemoryUsage_used		
		NonHeapMemoryUsage_committed		
		NonHeapMemoryUsage_init		
		NonHeapMemoryUsage_max		
		NonHeapMemoryUsage_used		
Java fd		OpenFileDescriptorCount		Search Store, Index Store

Table 3-1 Collected Performance Data (continued)

Type	Instance	Matrix	Description	Role
Non Java Application processes	ps_count	processes	Total number of processes (including child) forked for particular program.	Analytics Store, JSON Store, Cache, RabbitMQ
		threads	Total number of threads created for particular program.	
	ps_code			Analytics Store, JSON Store, Cache
	ps_data			Analytics Store, JSON Store, Cache
	ps_rss			Analytics Store, JSON Store, Cache
	ps_stacksize			Analytics Store, JSON Store, Cache
	ps_vm			Analytics Store, JSON Store, Cache
	ps_cputime	syst		Analytics Store, JSON Store, Cache
		user		
	ps_disk_octets	read		Analytics Store, JSON Store, Cache
		write		
	ps_disk_ops	read		Analytics Store, JSON Store, Cache
		write		
	ps_pagefaults	majfit		Analytics Store, JSON Store, Cache
		minfit		

Table 3-1 *Collected Performance Data (continued)*

Type	Instance	Matrix	Description	Role
MongoDB		cache_misses		Analytics Store, JSON Store
		connections		
		page_fault		
		lock_ratio%		
	flushes	flushes		
		flushes_avg_ms		
	memory	mapped		
		resident		
		virtual		
	network	bytesin		
		bytesout		
	oplogs	diffimesec		
		storagesizemb		
		usedsizemb		
	replication	health		
		optimelagsec		
		state		
	total_operations	command		
		delete		
		getmore		
		insert		
		query		
		update		
MongoDB databases	quad, recommendation	collections		
		indexes		
		num_extents		
		object_count		
		data file_size		
		index file_size		
		storage file_size		
Tomcat		activeSessions		App Server, Worker
		expiredSessions		
		processExpiresFrequency		
		processingTime		
		rejectedSessions		
		sessionAverageAliveTimes		
		sessionCounter		
		sessionCreateRate		
		sessionExpireRate		

Table 3-1 *Collected Performance Data (continued)*

Type	Instance	Matrix	Description	Role
RabbitMQ	Queue: Activity, Analytics, EMailDigest, Migrate, Polling, Scheduler	consumers		Message Queue
		memory		
		messages		
		messages_ready		
		messages_acknowledged		
		node		
	Server	fd_total		Message Queue
		fd_used		
		mem_limit		
		mem_used		
		proc_total		
		proc_used		
		sockets_total		
		sockets_used		
		uptime		
ActiveMQ Broker	TotalEnqueueCount			Message Queue
	TotalDequeueCount			
	TotalConsumerCount			
	TotalMessageCount			
	MemoryLimit			
	MemoryPercentUsage			
	StoreLimit			
	StorePercentUsage			
ActiveMQ Queue	QueueSize			Message Queue
	EnqueueCount			
	DequeueCount			
	ConsumerCount			
	DispatchCount			
	ExpiredCount			
	InFlightCount			
	CursorMemoryUsage			
	CursorPercentUsage			
	MemoryLimit			
	MemoryPercentUsage			

Monitored Health Metrics

This section summarizes the resources that are monitored by monit to ensure good health of the system. Monit automatically takes corrective action if a process stops or becomes unresponsive. A syslog message is generated on alert and when corrective action is taken. Monit checks are only done on Enabled applications.

This data can be accessed in several ways:

- From the Director UI > System > Health.
- Through the WebEx Social API.

Table 3-2 Monitored Health Metrics

CheckName/ Filename	Type	Checks	Action	Role
jms-message-queue/ process_activemq	Process	pid	Restart	Message Queue
		cpu > 98% for 5 poll	Syslog Err Msg	
analyticsstore/ process_analyticsstore	Process	pid	Restart	Analytic Store
		tcp on port 27001 for 1 poll	Syslog Err Msg	
analyticsstore/ process_analyticsstore ¹	Process	pid	Restart	Director
		tcp on port 27001 for 1 poll	Syslog Err Msg	
		cpu > 98% for 5 poll	Syslog Err Msg	
cache/ process_cache	Process	pid	Restart	Cache
		Built-in monit protocol check for memcache on port 11211 for 1 poll	Syslog Err Msg	
		cpu > 98% for 5 poll	Syslog Err Msg	
carbon/ process_carbon	Process	pid	Restart	Director
		cpu > 25% for 5 poll	Syslog Err Msg	
cmanager/ process_cmanager	Process	pid	Restart	WebEx Social
		cpu > 98% for 5 poll	Syslog Err Msg	
collectd/ process_collectd	Process	pid	Restart	All
		cpu > 25% for 5 poll	Syslog Err Msg	
director-web/ process_cps	Process	pid	Restart	Director
		cpu > 98% for 5 poll	Syslog Err Msg	
	Disk Space	/opt > 85% for 5 poll	Purge /opt/logs/*. Except today's log	
cron/ process_cron	Process	pid	Restart	All
httpd/ process_httpd	Process	pid	Restart	Director, WebEx Social, Worker
indexstore/ process_indexstore	Process	pid	Restart	Index Store
		cpu > 98% for 5 poll	Syslog Err Msg	
jsonstore/ process_jsonstore	Process	pid	Restart	JSON Store
		tcp on port 27000 for 1 poll	Syslog Err Msg	
		cpu > 98% for 5 poll	Syslog Err Msg	
jsonstore/ process_jsonstore ²	Process	pid	Restart	Director
		tcp on port 27000 for 1 poll	Syslog Err Msg	
		cpu > 98% for 5 poll	Syslog Err Msg	
nagios/ process_nagios	Process	pid	Restart	Director
		cpu > 25% for 5 poll	Syslog Err Msg	
ntpd/ process_ntpd	Process	pid	Restart	All
		cpu > 25% for 5 poll	Syslog Err Msg	

Table 3-2 Monitored Health Metrics (continued)

CheckName/ Filename	Type	Checks	Action	Role
notifier/ process_openfire	Process	pid	Restart	Notifier
		cpu > 98% for 5 poll	Syslog Err Msg	
postfix/ process_postfix ³	Process	pid	Restart	Director, Worker
		cpu > 40% for 2 poll	Syslog Err Msg	
		cpu > 60% for 5 poll	Restart	
		Built-in monit protocol check for SMTP for 1 poll	Syslog Err Msg	
		Children > 2000	Syslog Err Msg	
		Memory > 2GB for 2 poll	Restart	
puppet/ process_puppet	Process	pid	Restart	All
		cpu > 98% for 5 poll	Syslog Err Msg	
puppetmaster/ process_puppetmaster	Process	pid	Restart	Director
		tcp on port 8140 for 1 poll	Syslog Err Msg	
		cpu > 98% for 5 poll	Syslog Err Msg	
quad/ process_quad	Process	pid	Restart	WebEx Social
		cpu > 98% for 5 poll	Syslog Err Msg	
message-queue/ process_rabbitmq	Process	pid	Restart	Message Queue
		cpu > 98% for 5 poll	Syslog Err Msg	
rsyslog/ process_rsyslog	Process	pid	Restart	All
		tcp on port 514 for 1 poll	Syslog Err Msg	Director
		cpu > 50% for 5 poll	Syslog Err Msg	All
saltmaster/ process_saltmaster	Process	pid	Restart	Director
		tcp on port 4506 for 1 poll	Syslog Err Msg	
		cpu > 98% for 5 poll	Syslog Err Msg	
saltminion/ process_saltminion	Process	pid	Restart	All
		cpu > 98% for 5 poll	Syslog Err Msg	
search/ process_searchstore	Process	pid	Restart	Search Store
		cpu > 98% for 5 poll	Syslog Err Msg	
sshd/ process_sshd	Process	pid	Restart	All
		Built-in monit protocol check for ssh on port 22 for 1 poll	Syslog Err Msg	
		cpu > 25% for 5 poll	Syslog Err Msg	
worker/ process_worker	Process	pid	Restart	Worker
		cpu > 98% for 5 poll	Syslog Err Msg	
oracle/ program_oracle ⁴	Program (script)	script return value; for 10 polls	Restart	RDBMS Store, Graph Store
integrity/ program_integrity	Program (script)	script return value;	Syslog Err Msg	All
Disk usage check ⁵	/opt	> 85%	Nagios Warning	All
	/opt	> 95%	Nagios Alert	
	/boot	> 99%	Nagios Alert	
	/root	> 99%	Nagios Alert	

1. Arbiter check available only where there are multiple Json/Analytics VMs.
2. Arbiter check available only where there are multiple Json/Analytics VMs.
3. Postfix service monitored only when maildomain/external host and external SMTP port are provisioned.

4. The check is done using “/etc/init.d/dbora status”. Restarting is done using “/etc/init.d/dbora cond_start”. Only services that are not running (Enterprise Manager, Database etc) are started. Checks are not made during database installation.
5. The disk utilization check uses performance statistics as collected by collectd.

