Software INTERNATIONAL Concepts

World Class Managed
Services for OpenVMS

Best Practices for OpenVMS System and Rdb Management for the 21st Century

Bryan Holland

Software Concepts International, LLC
402 Amherst Street, Suite 300
Nashua, NH 03063, USA

Phone: 603-879-9022
e-mail: holland@sciinc.com
www.sciinc.com

Agenda

- Introduction
- Current state of OpenVMS Systems
- Possible Problems
- Recommendations

Agenda

- Introduction
- Current state of OpenVMS Systems
- Possible Problems
- Recommendations

About
Software Concepts International

- Located in Nashua, NH (USA)
- 23+ years in business supporting OpenVMS
- An international reputation
  - A leading provider of remote managed DBA services for the Rdb and DBMS databases
  - A leading provider of remote managed services for OpenVMS systems
- Proven track record
  - Actively managing 100s of databases and dozens of systems and configurations
  - Remote DBA service since 1995 (still supporting many of the same sites)

Session Objectives

- Discuss OpenVMS system management issues for today’s OpenVMS environments
- Describe methods and options for VMS System Management support into the future.

Non – Objectives

- We do not plan to discuss *detailed* System Management solutions
  - No code samples
- No philosophical discussions of OpenVMS system management

Agenda

- Introduction
- Current State of OpenVMS Systems
- Possible Problems
- Recommendations

OpenVMS Today

OpenVMS provides the core IT infrastructure for:

- Mobile phone billing systems scaling to millions of users
- Major futures and derivative exchanges worldwide
- Majority of automated lottery systems
- Many of the world’s most demanding Government environments requiring security and availability
- Manufacturing from CPU chips to automobiles

OpenVMS – Rock Solid

- Highly Reliable, Available, Secure
- VMS is extremely resilient
- Highly disaster tolerant
- Security built in from the ground up
- Uptime often measured in years
- Hardware is equally reliable

almos

These systems never go down!

t^

Bet Your Business

- Most OpenVMS systems are mission or business critical
  - Application availability is essential to keeping business processes running
  - Unplanned system or application downtime is usually a severe problem

- What is the impact to your business if your OpenVMS system is not available?

Current State

- Many systems were put into service 15-20 years ago (1990s)
- The VMS experts have moved on, retired, or are no longer available
- Systems often receive a bare minimum of updates and maintenance
- Many systems sit in a corner and run.

It’s easy to become complacent…

Complacency is Dangerous

- Has the reliability made you complacent?
- Are you able to recognize and prevent a problem *before* it happens?

Complacency is Dangerous (cont)

- Are your systems generating logs or reports?
  - – Is anyone monitoring this output?
  - – Will it be obvious when a real problem occurs?
    - Are “known problems” being ignored?
- Are changes reflected in startup files?
  - – When you reboot, will the environment start correctly

Complacency is Dangerous (cont)

- Do you maintain historical data?
  - Baseline configuration information
  - Performance data
  - System resource and parameter data

Complacency is Dangerous (cont)

- When something does go wrong
  - Do you have everything you need to repair, recover, restart? (You need these things BEFORE the failure)
  - Will you know for certain that everything has been restored to the prior running state?

Complacency is Bad (cont)

- Who is your VMS system manager?
  - – OpenVMS expert?
    - How accessible? Available 24x7?
  - – UNIX/Linux/Windows experts?
    - These people are OK, as long as nothing abnormal occurs

Understand Your Risks

- Do you *really* know your risks?

- If the application is suddenly unavailable…
  - Will your customer/users accept “But, it’s been available non-stop for more than a year”
  - Will you have the resources (tools, procedures, valid backups, knowledgeable people) to restore the application?

Lights Out Management

- Lights Out is the standard
  - Assumes tools and processes in place for Lights Out management

- Are your VMS systems Lights Out managed?
  - Or did someone *just turn out the lights and forget about it*?

Rock Solid?

Rock Solid?
Things Can Go Wrong

- Do you have “peace of mind” that you have done everything you can to ***prevent*** a failure?
  - It’s not just catastrophic incidents to worry about – small things can lead to outages

Rock Solid?
Things Can Go Wrong

- Do you have “peace of mind” that you are prepared to quickly ***recover*** from a failure?
  - – Recover completely, quickly and confidently?
  - – With no loss of data?

Agenda

- Introduction
- Current State of OpenVMS Systems
- Possible Problems
- Recommendations

OpenVMS is resilient,
But,…
it’s not self maintaining

What Can Go Wrong?

- Many things…
  - – Simple or Complex
  - – H/W, S/W, M/W, other systems
  - – Environmental
  - – Users
  - – Security
- It may be one or many of the above that cause application unavailability

Disk Space

- **Problem**: Running out of disk space
- **Action**:
  - – Monitor disk utilization regularly
  - – Warn when thresholds are crossed
    - Use multiple levels – e.g. Warning, Severe, Critical
  - – Preserve statistical data over time –
    - Know what is normal usage for each disk
    - Predict problems

Disk Utilization charts

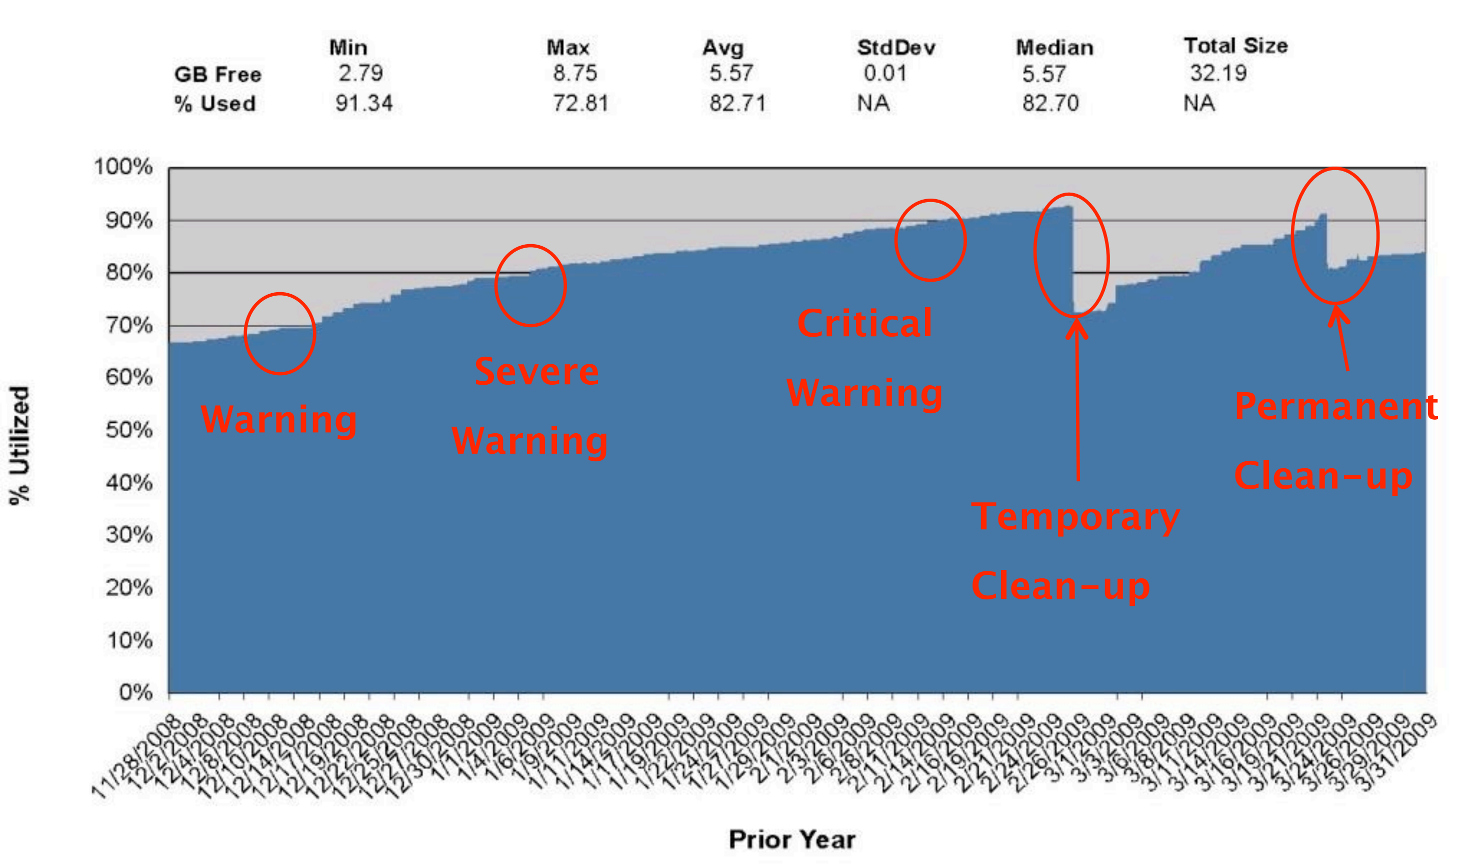

	Min	Max	Avg	StdDev	Median	Total Size
GB Free	2.79	8.75	5.57	0.01	5.57	32.19
% Used	91.34	72.81	82.71	NA	82.70	NA

Warning

Severe Warning

Critical Warning

Temporary Clean-up

Permanent Clean-up

Disk Utilization Charts

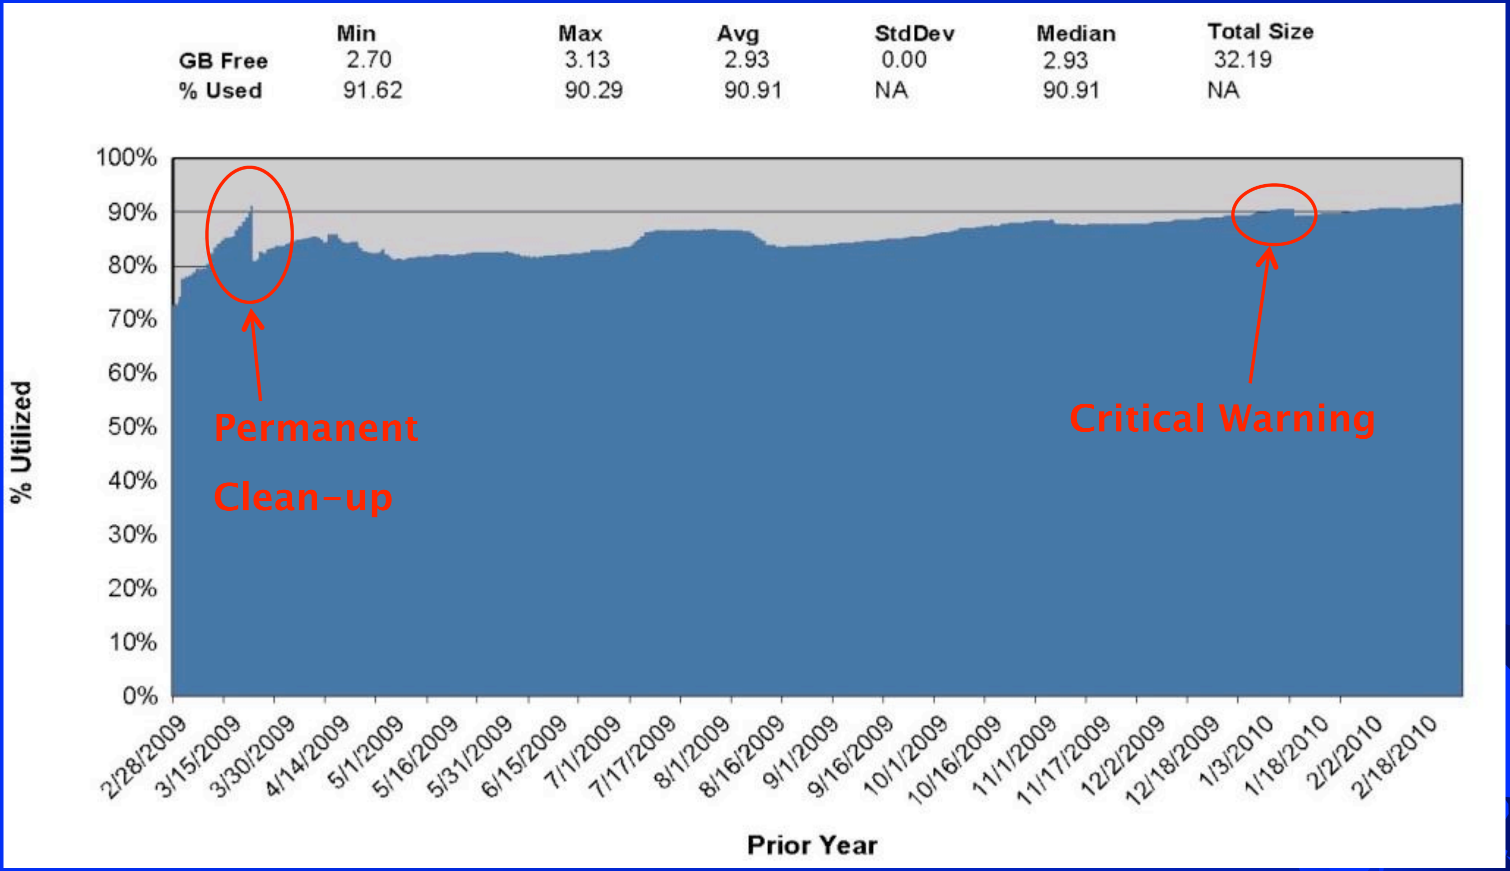

	Min	Max	Avg	StdDev	Median	Total Size
GB Free	2.70	3.13	2.93	0.00	2.93	32.19
% Used	91.62	90.29	90.91	NA	90.91	NA

Permanent Clean-up

Critical Warning

File System

- **Problem:** File System issues can lead to application problems/interruption
  - Excessive number of files
    - I.E. application not cleaning up after itself.
  - High file version number
  - Directories too big

File System (cont)

- **Action:** Scan directories/device regularly for problem files
  - Use multiple levels – e.g. Warning, Severe, Critical
  - Regular scheduled cleanup, if appropriate
  - Consider version limits

Disk Device

- **Problem**: Disk device errors could be indicative of coming hardware problems

- **Actions**: Monitor and report, take steps to avoid loss of data.
  - Monitor regularly, report errors
  - Repair or replace disk

Disk/File System - Corruption

- **Problem**: File System inconsistencies, lost files, file system errors, quota inconsistencies
  - – Could result in data corruption or data unavailability
- **Action**: ANALYZE/DISK_STRUCT
  - – Recommend running regularly
  - – Report errors found
  - – For some errors, consider automatic "/ REPAIR"

Disk/File System - Corruption (cont)

- Recommend the following errors be fixed automatically:
  - – Backlink errors
  - – Bitmap using more space
  - – Disk quota errors
  - – Files marked delete

- Reporting should include a list of lost files for examination

Disk/File System - Performance

- **Problem:** Performance not optimal
  - – Excessive file read/write activity
  - – Particularly for heavily used index files or databases
  - – Disk path not optimal (e.g. MSCP vs. FC)
- **Action:**
  - – Monitor XFC cache
    - Check for ineffective utilization
    - Hot files

Disk/File System – Performance (cont)

- **Problem**: Performance problems (cont)
- **Action**:
  - – For index files – anal/rms, convert with tuned FDLs
  - – Monitor for fragmentation
    - Run defragmenter, as needed if you have one
    - Image/backup if no defragmenter

STARTUP Monitoring

- **Problem**: Errors during boot process
  - Could result in errors long after boot is completed
  - Could impact production application availability

- Action: Enable Startup logging and scan for errors

Startup Logging

- SYSGEN **`STARTUP_P2 = CDV`**
  - “V” – creates *a lot* of logging
  - Update MODPARAMS.DAT
- Upon completion of boot, scan file STARTUP.LOG for errors
  - “-E-”, “-F-”, “-W-”
  - Report on and repair any errors found

Startup Logging

- Don’t forget to log and review:
  - Spawned processes
  - Submitted startup routines
  - Any startup performed outside the VMS STARTUP phase.
- All of these should be automatically scanned and errors (or new errors) reported.

Shutdown

- **Problem**: Errors during shutdown can cause problems upon or after reboot

- **Actions:**
  - Use shutdown procedures
    - Cleanly shutdown applications and databases
  - Scan logs on startup so you know if there were problems

System Logs

- **Problem**: System Logs get too big, difficult to search and manage
  - Accounting, Operator, Security
- **Actions:** Periodically create new logs
  - Recommend automated periodic job
  - Archive old logs with date embedded file names
  - Retain per local requirements

A Word About Security…

- Next few slides will NOT tell you how to completely secure your system
  - That could be a full day seminar.
- Government regulations or industry standards may dictate specific security requirements

Security Monitoring

- **Problem**: Security can be breached

- **Actions:** Enable Security Auditing
  - – Periodically check the audit and alarm settings for changes.
  - – Real-time alerting for security Alarms for specific events of interest at your site, e.g.:
    - ACL auditing of critical files
    - AUTHORIZE auditing
    - Break-In

User Management

- **Problem:** Changes to User Accounts may create security risks

- **Actions:** Periodic scans of SYSUAF
  - Last login times >X days (privileged)
  - Last login times >Y days (non-priv’d)
  - Large number of login failures

User Management

- **Problem:** Changes to User Accounts

- **Actions:** Monitor changes in SYSUAF
  - New accounts, deleted accounts, privileges, proxy changes
  - Consider enabling Authorize Security Audits
    - Monitor security audit log file

Installed Images

- **Problem**: Unexpected changes to installed images could be
  - Hidden security risks
  - Unapproved application changes
- Action: Periodically monitor installed images
  - Use INSTALL to look for changes in known file list

T4/Performance Data

- **Problem:** Lack of data for performance problem investigation
- **Actions:**
  - Setup T4 to collect CPU usage, I/O activity, and cluster traffic (if applicable)
  - Roll up CSV files into appropriate intervals for reporting
  - Graph results for ease of analysis
  - Maintain historical data for comparison

Critical Files

- **Problem:** Important events go unnoticed
- **Actions**: Monitor critical files
  - **Critical Files are:**
    - Log files – presence or lack of presence is significant
      - Expect to see a file at some frequency – report when you don’t see it
      - Expect to see a file only when there is problem – report when seen
    - Open log files may take some creativity to monitor

Critical Files

- **Problem**: A file was changed when it shouldn’t have been
  - Content has changed
- **Actions**:
  - Identify files that should not change
  - Consider setting ACL Security Audits/Alarm
  - Scan periodically – compare states and report changes
  - Three states: current, last, known good

System Parameters

- **Problem**: Unexpected changes to SYSGEN parameters
  - Changes need to be approved and correctly implemented.
- **Action**: Monitor SYSGEN parameters
  - Dump parameters to a file (Active and Current)
  - Compare states, report on changes
  - Current, last, last known good.

System Parameters (cont)

- Could also set a security ACE on parameter file
  - SYS$SYSROOT:[SYSEXE]*VMSSYS.PAR
  - Trigger warning/alert on security alarm
  - Follow up with manual inspection

System Resources

- **Problem:** Key system resources get used beyond system limits
  - System can hang/crash
  - Applications may fail
  - Performance degrades

System Resources (cont)

- **Action**: Monitor resource vs. corresponding parameter and alert as appropriate
  - May be several heuristics that apply
    - Percentage changed and percentage consumed
      - E.g. NPAGEVIR, PAGEDYN
    - Amount changed and Amount consumed
      - E.g.   Process Count
    - Change rate

System Resources (cont)

- Run Autogen periodically to create report
  - Review report, take action if necessary
  - Preserve reports for historical reference
- Run Autogen periodically to allow the system to tune itself

Logical Names

- **Problem:** Shared logical changes (SYSTEM, GROUP, etc)
  - Lack of persistence across reboots
  - Unexpected changes

- **Actions:** Establish baseline and monitor
  - Baseline – immediately after reboot
  - Monitor known critical logicals, report changes
    - current, last, known good

Environment changes

- **Problem**: Some environment changes don’t survive reboots

- **Actions:** Implement environmental changes via established, documented processes.
  - Monitor critical resources for changes
    - Files, queues, logical names
  - Insure changes survive reboots

Licenses

- **Problem**: License expires, critical product stops working

- **Action**: Monitor the license database
  - Look for Termination Dates nearing expiration
  - Check other non-LMF licenses
    - E.g. BEA MessageQ – a text file which can be scanned

Networks

- **Problem**: Network *Interfaces* Fail
  - Possibly fail unknowingly
- **Action**: If available…
  - Set up TCP/IP failSAFE or LAN Failover
  - Monitor and warn if failover has occurred
    - TCPIP$SYFAILSAFE.COM can trigger warnings
    - LANCP SHO DEV LL /CHAR

Networks

- **Problem**: Unexpected Network Service Changes

- **Action**: Monitor network services
  - **TCPIP SHOW SERVICE**
  - **TCPIP SHOW SERVICE xxxx /FULL**

Queues

- **Problem:** Queues not available for processing
- **Actions:** Monitor queues periodically
  - Watch for unexpected queue states (stalled, stopped)
  - Watch for stuck jobs – queue busy with same job longer than expected
  - Automatically restart, when appropriate
  - Also monitor queue attributes for changes

Redundancy

- **Problem**: Expected redundancy not available
  - – Redundancy can be so transparent that you don’t notice when a component fails
- **Actions**:
  - Monitor Raid set members
  - Monitor Fibre Channel paths
  - Network Interfaces

Patches

- Problem: Important patches not installed
  - 3/23/2010 – HP release HPSBOV02497, Security risk in NTP. Did you know?
- Actions:
  - Sign up with HP for patch notifications
  - Maintain list of outstanding patches
  - Apply patches when appropriate or necessary

Rdb Reliability/Availability

- Physical separation of database files (.rdb, .rda, .ruj) from
  Recovery files (.rbf, .aij)
- Enable (multiple-”circular”) AIJs
  (Place on shadowed/mirrored device)
- Create, test and execute recovery strategy

Rdb Recovery

- AIJ backup
- DB backup
- Hot-standby (real-time synchronization)
- Warm-spare (periodic synchronization)
- Only works if running successfully (must monitor)

Early Rdb Problem Detection

- Real-time scan of Rdb Monitor log file (s) [each node/version]
- Real-time scan of OPCOM messages for Rdb messages (AIJ problems)
- Define RDM$BUGCHECK_DIR – scan/ review bugchecks.
- Enable and monitor Rdb server logs (ABS, DBR, LCS, LRS, RCS)

Rdb problem detection

- Perform regular Verifies (internal integrity) of the database
  - Use a restored copy or hot-standby
- Processes with Old TSNs
  - Snapshot growth

Rdb Security

- Set RMU/database/table protections
- Enable (and monitor) Rdb auditing
- Real-time monitoring of Rdb “alerts”
- Monitor for unanticipated changes:
  - RMU/DUMP
  - RMU/EXTRACT/ITEM=ALL (Schema/ security changes)
  - RMU/SHOW PRIVILEGES (RMU privileges)

Rdb Performance Management

- Collect, monitor and analyze historical information regarding area utilization (RMU/ANALYZE)
- Collect, monitor and analyze historical run-time statistics data (RMU/SHOW STAT)
  - Maintain historical performance data
  - TLVIZ

Agenda

- Introduction
- Current State of OpenVMS Systems
- Possible Problems
- Recommendations

Reactive Support

- Model used by many
  - Something breaks, then
  - Heroic actions ***may*** result in a fix
    - Resolving problems under pressure never optimal
    - Solution probably not repeatable either
    - Root Cause Analysis?

Reactive Support (cont)

- May implement change to avoid problem again
  - Seldom is holistic approach taken to avoid all similar situations

Reactive = Too Late!

- Goals should be:
  - Avoid problems
    - Zero unplanned down time
  - Efficiently restore services
  - Prevent future occurrences

Proactive = Just Right!

Proactive Support

- Establish a baseline
  - Critical to know what “normal” is
  - Should keep track of:
    - “Known Good” - This is the baseline
    - “Last Known” – this is what was seen on the last scan – important to know if things are continuing to change

Proactive Support

- Repeatable and Consistent Processes
  - Monitor for changes
  - Watch Trends
  - Correct problems

- Proactive is NOT “periodic login” to “review logs”

Proactive Support

- Problems/Incidents
  - Things that are not supposed to happen
  - Things that should happen, but don’t
  - Thresholds exceeded
  - Trends – possible problems coming
- Automated notification – ideally with required acknowledgement
  - E.g. phone calls that keep trying until answered

Proactive Support

- Processes should integrate with a persistent Trouble Ticket System
  - Persistent does not mean “send e-mail” (aka fire and forget)
  - Persistent means there is an ongoing tracking and accounting of progress to completion.
  - Full accountability to getting problem fixed.

Proactive Support

- Trouble Ticketing System
  - All reportable events must be trapped and integrated with trouble-ticket reporting system
  - Critical events must alert appropriate support staff immediately & integrate with Voice Response Systems
  - A checklist is required to validate the successful completion of scheduled tasks
  - Reporting is required for missing/late tasks.

Proactive Support

- Trouble Ticketing System (cont)
  - All incidents should result in a ticket being created
  - Forces accountability – problems can't be forgotten or ignored
    - Forces someone to take action
  - Monitoring tools should automatically create (and close) tickets.
    - Logging an error or just sending an e-mail isn't sufficient

Proactive Support

- Holistic approach
  - Identify the problem
  - Document and Track progress
  - Fix
  - Root cause analysis
  - Prevent

Conclusion

Managed Systems

Support Systems

Staff

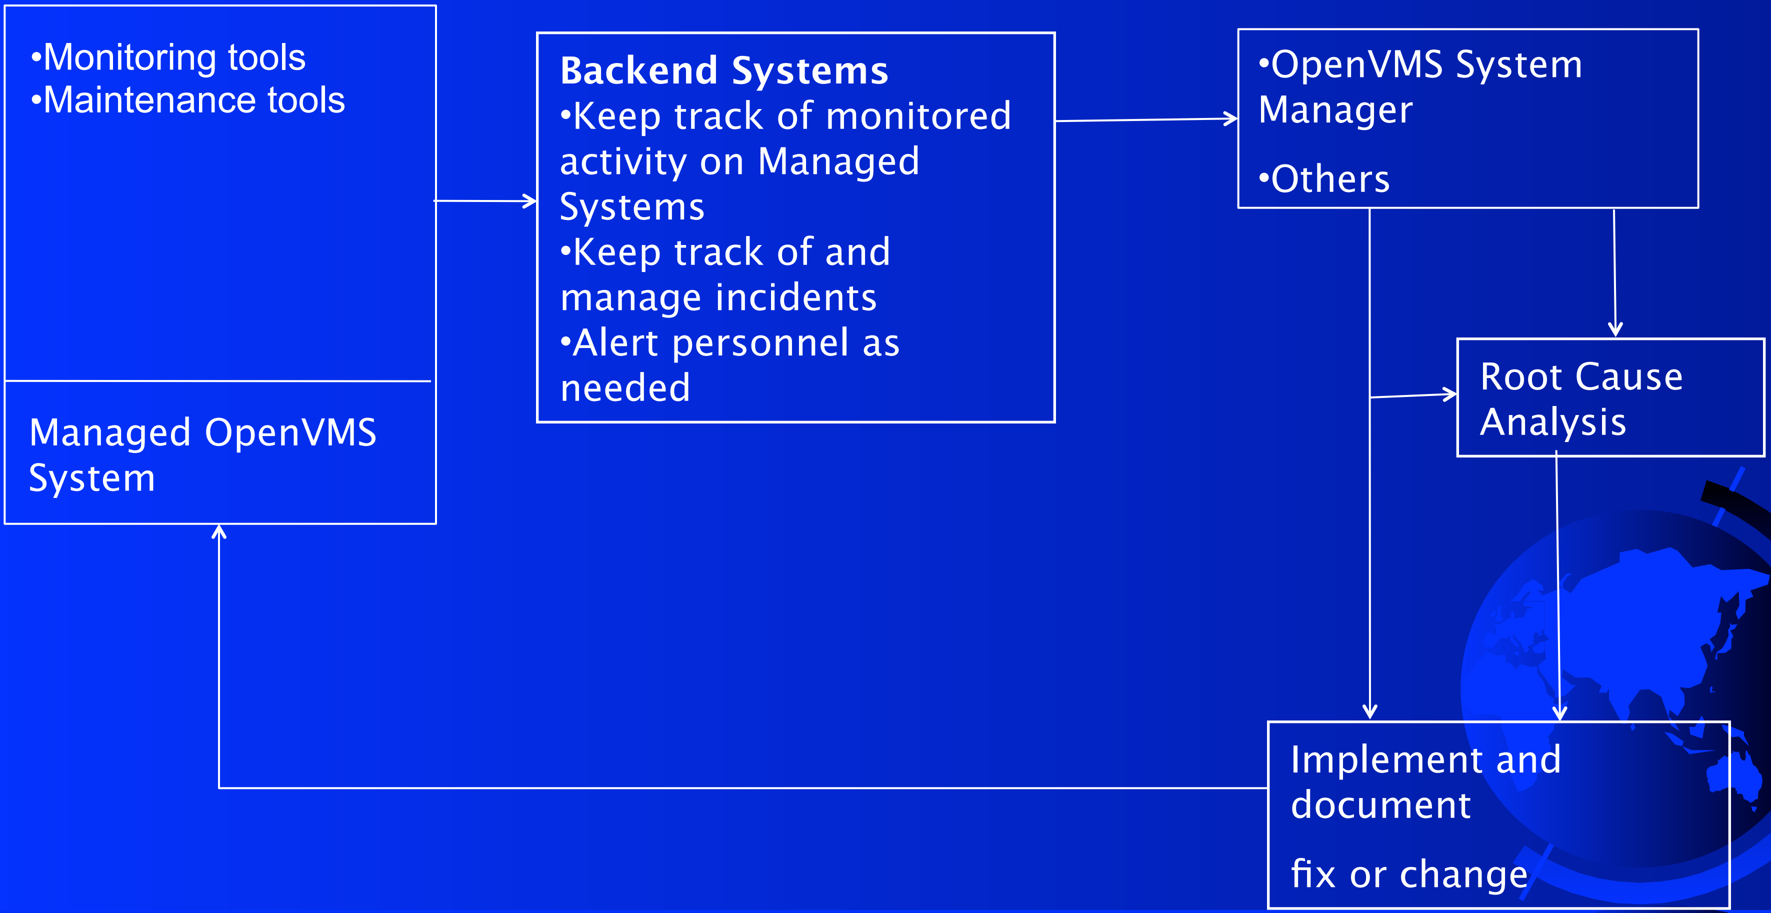

- Monitoring tools
- Maintenance tools

Managed OpenVMS System

Backend Systems

- Keep track of monitored activity on Managed Systems
- Keep track of and manage incidents
- Alert personnel as needed

- OpenVMS System Manager
- Others

Root Cause Analysis

Implement and document

fix or change

Conclusion

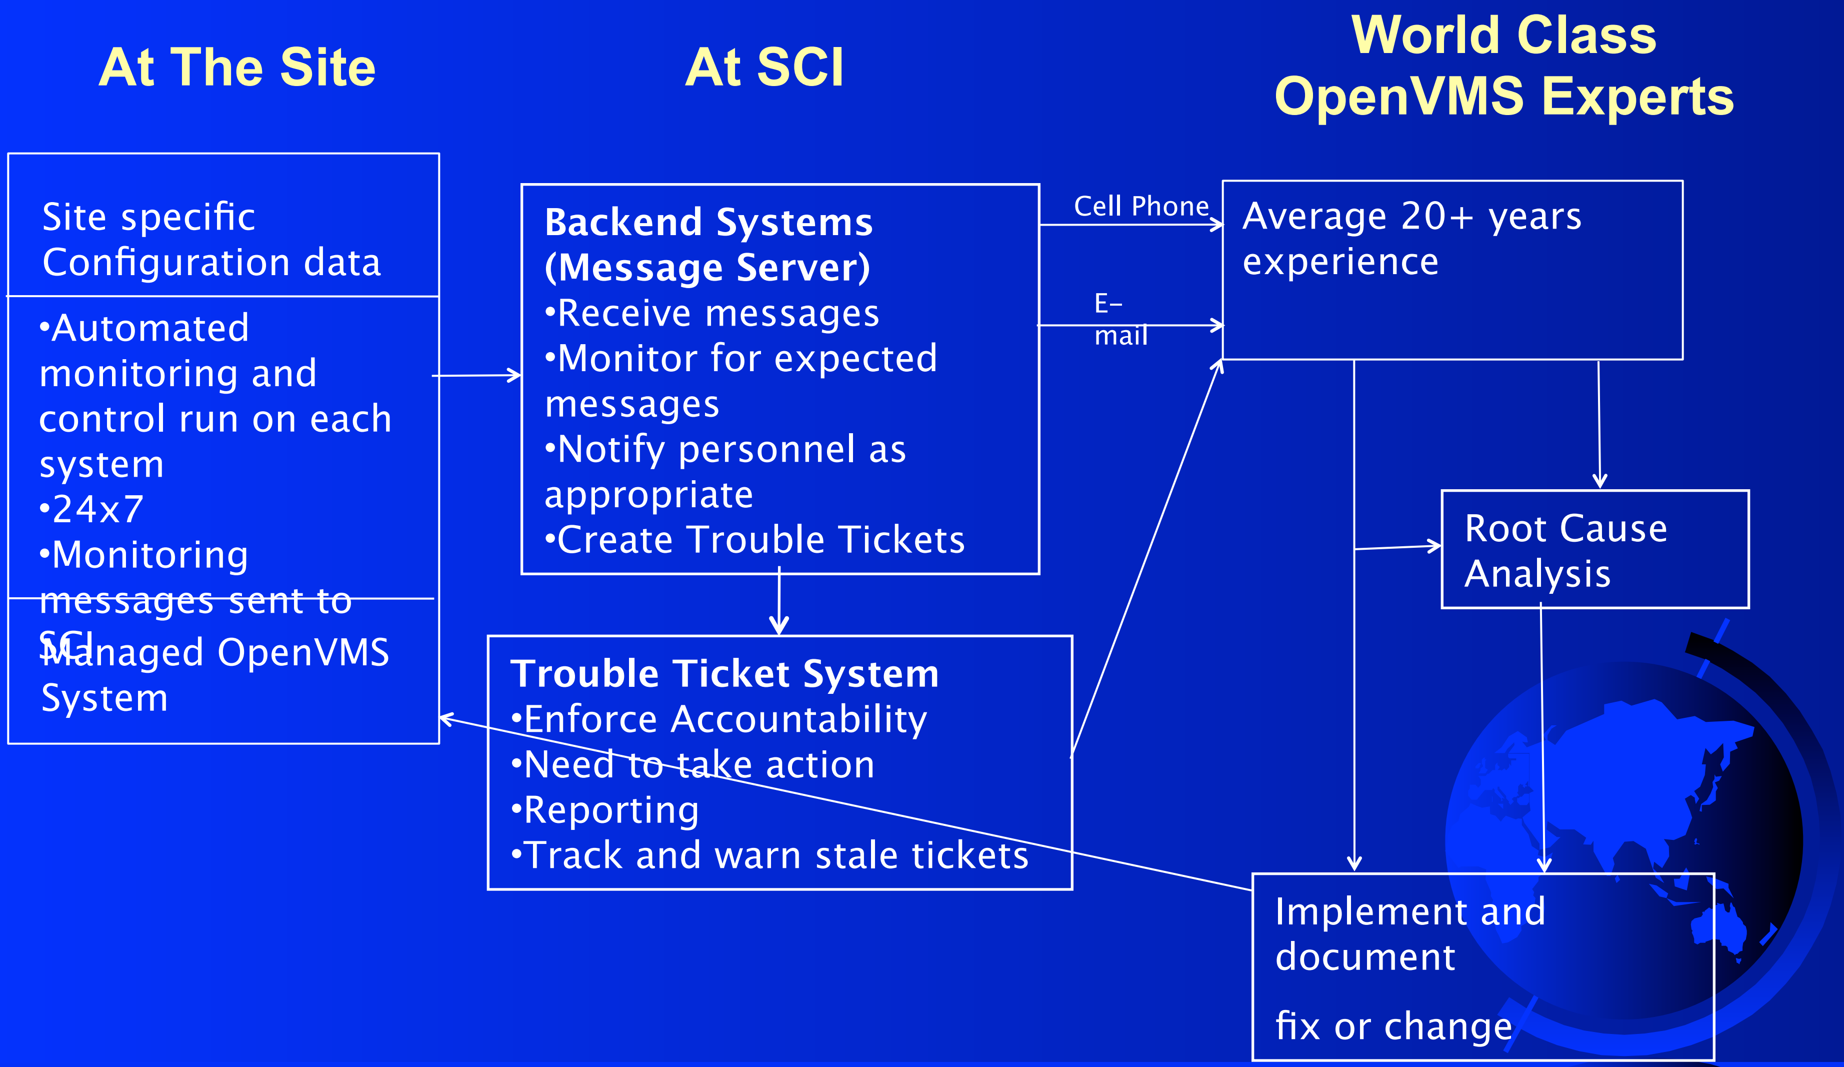

At The Site

Site specific Configuration data

- Automated monitoring and control run on each system
- 24x7
- Monitoring messages sent to SCI

Managed OpenVMS System

At SCI

Backend Systems (Message Server)

- Receive messages
- Monitor for expected messages
- Notify personnel as appropriate
- Create Trouble Tickets

Trouble Ticket System

- Enforce Accountability
- Need to take action
- Reporting
- Track and warn stale tickets

World Class OpenVMS Experts

Average 20+ years experience

Root Cause Analysis

Implement and document

fix or change

Some Closing Thoughts

- Previous Slide:
  - This is what you should strive to build
- What business are you in?
- Focus on your business's core competencies

Questions?