



Connect Germany

G T U G
G O I N T E R N A T I O N A L

Abstract Workshop III
14 May 2018, 14.00 h

XDelta Limited

Speaker

Colin Butcher, in collaboration with Keith Parris (HPE) and Nic Clews (DXC)

Speaker Biography

Colin is a consulting engineer who specialises in leading the design and implementation of mission-critical and disaster-tolerant systems. He started XDelta Limited in 1996 and works with a wide range of OpenVMS based systems and their surrounding infrastructure. Colin has been responsible for the design and implementation of systems engineering projects in a wide range of industries such as aerospace, power generation, finance, healthcare and transport.

Colin has been involved with the world-wide OpenVMS community for over 30 years and is a well-known speaker with specific interests in availability, performance, security, networking and storage. His background is in large scale systems engineering, including software design and development.

Title

Best practices for OpenVMS systems and clusters.

Overview

This workshop will provide information gained by practical experience with the design, build, deployment and trouble-shooting of OpenVMS systems and clusters, with emphasis on high-availability multi-site OpenVMS clusters. The content has been contributed to by several people, who have kindly made the effort to share their knowledge and experience. We hope that it will prove useful to increase your level of understanding and enable you to make your own decisions based on a good understanding of the wide range of topics involved when working on such systems.



Connect Germany

G T U G
G O I N T E R N A T I O N A L

Abstract Workshop III
14 May 2018, 14.00 h

Abstract I

The workshop will include the following subject areas:

- OpenVMS boot, startup and shutdown sequences.

Understanding the boot sequence is important, especially on Integrity Servers, where the view from EFI shell is very different to that provided by OpenVMS. Using "dump off system disk" (DOSD) has specific requirements. Setting up OpenVMS systems and clusters requires writing DCL to handle the startup and shutdown of the complete system, including the applications. A consistent and modular approach is important in order to minimize the difficulty of trouble-shooting as changes are made over time.

- Making changes to operational systems with minimal disruption - firmware updates, operating system updates and hardware replacement.

Careful planning and representative testing are crucial. It is essential to verify the exact sequence of operations required and the approximate timings of each step. There are many pitfalls and traps waiting for the unwary. Working swiftly and accurately under time pressure can be extremely stressful, especially when unexpected events occur. Checklists are invaluable.

- Monitoring and alerting.

This is a key part of the operational regime. Without good monitoring of the end to end data flows and behavior, it can be extremely difficult to trouble-shoot problems. With high availability systems that have no single points of failure, monitoring and alerting is essential as the first failure of a component is "silent" in that nothing stops working from a system viewpoint. The second failure is the one that creates panic, when things suddenly stop working. Most operational problems are usually a combination of failures and events that happen to coincide in time to create a difficult and complex system failure.



Connect Germany

G T U G
G O I N T E R N A T I O N A L

Abstract Workshop III
14 May 2018, 14.00 h

Abstract II

- OpenVMS cluster design and implementation.

OpenVMS clusters can provide exceedingly high availability. The key to achieving this is to design the clusters carefully and plan ahead for change with minimal disruption. For example, there may be a change of hardware platform that requires different versions of the operating system and different system disks in order to boot, or it may be necessary to move data centres without disruption to service.

The surrounding storage and data network infrastructure is at least as important as the hardware platform.

Shadowing provides synchronous data replication in an OpenVMS environment. The maximum write transaction rate of the application is limited by the maximum achievable IO write rate to disk, which is determined by the worst-case write latency to the storage devices. In a multi-site cluster, the inter-site distance is the governing factor that determines the write latency to remote storage.

Data networking for the cluster interconnect is another key factor. The cluster interconnect is used for many things, key items being locking, mini-copy bitmaps and MSCP serving. Separation of traffic flows through the data network infrastructure can provide better control and management of the different traffic flows between the member nodes.

Topics to be covered include:

- Overall structure (system disks, system roots, page/swap/dump disks, common disk, data disks)
- Quorum and voting
- Storage connectivity with fibrechannel
- Network connectivity and separation of traffic flows
- Storage devices and shadowing
- Performance: latency and the effect on write IO rates in multi-site clusters, NUMA, locking
- Physical hardware layout (servers, storage, networks, rack space)

Come along and join in the discussion.