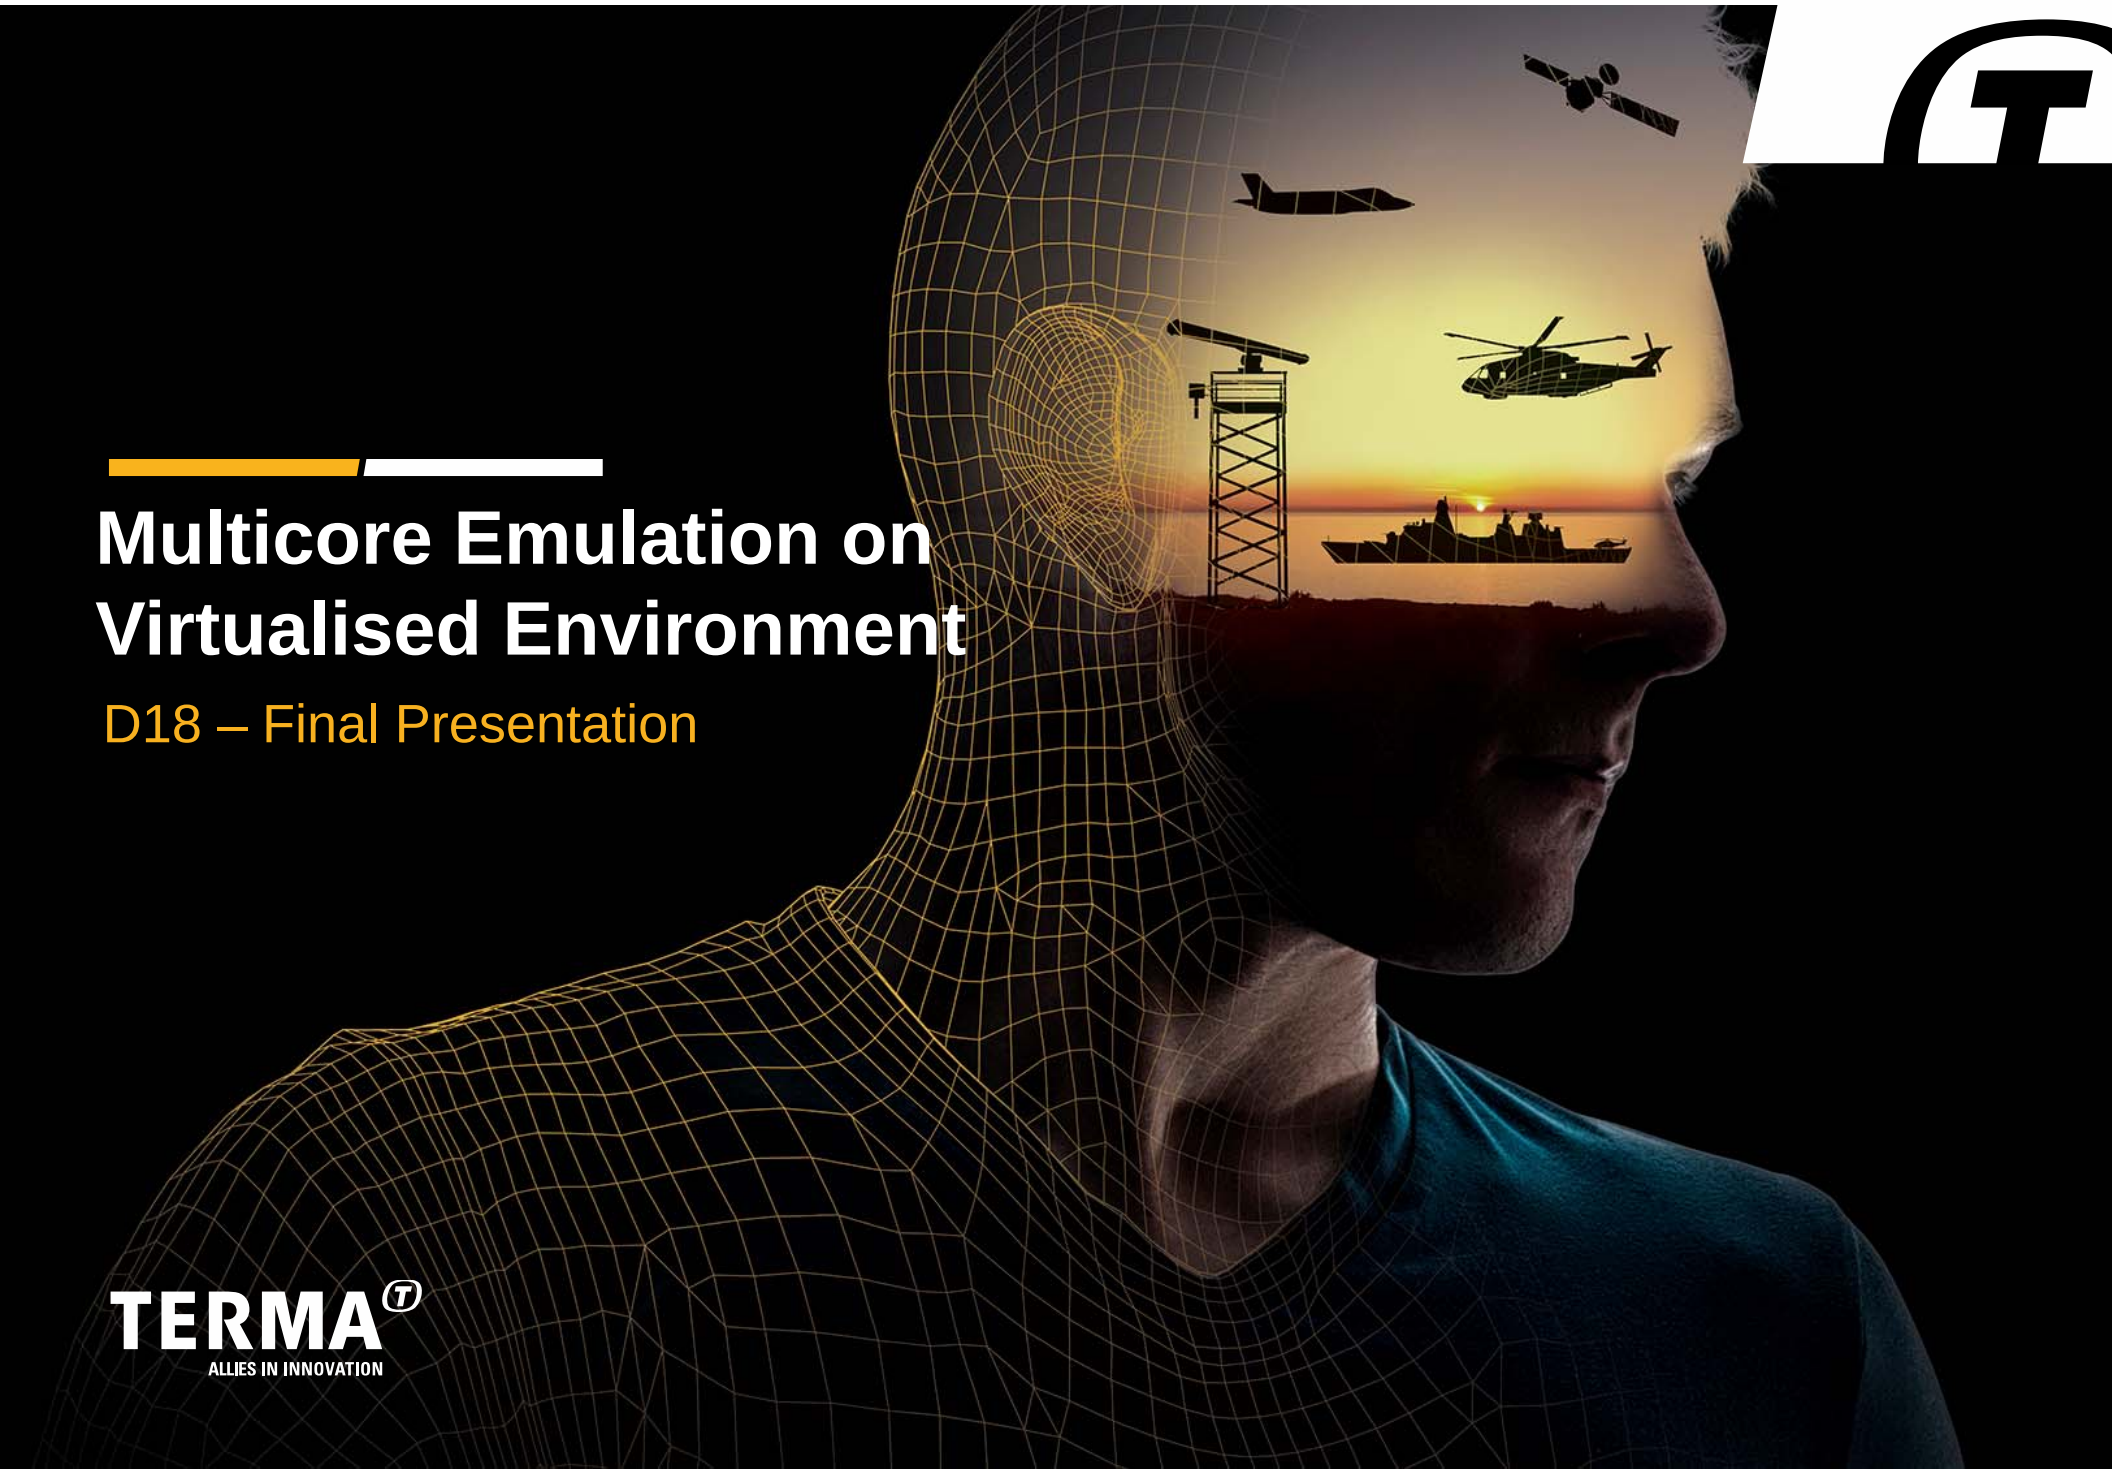

Multicore Emulation on Virtualised Environment

D18 – Final Presentation

TERMA

ALLIES IN INNOVATION

Agenda

- **Virtualisation and Real-time Simulation**
  - Use Cases
  - Virtualisation vs. Real-time
- **Test Suite**
  - Overview
  - Modularity and Extendibility
  - Test Cases, Time Reporting Server, etc.
  - Report Generation
- **Results**
- **Guidelines**

© The Terma Group 2017

Virtualisation - Quick Introduction (Xen)

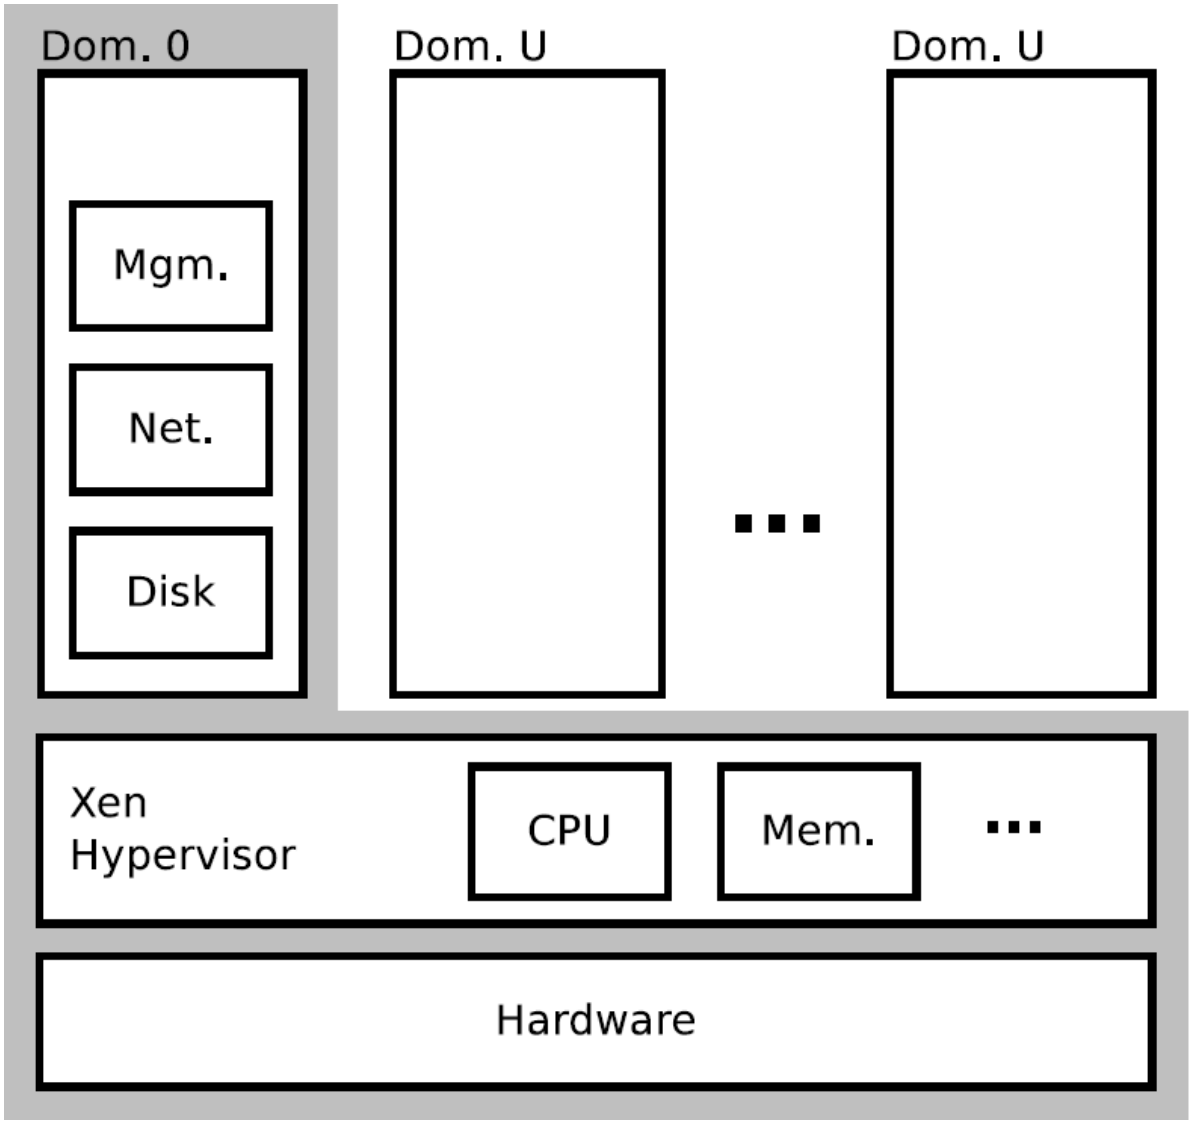

© The Terma Group 2017

Use Cases

- **IT Infrastructure Migration**
  - Physical -> Virtualized
- **Virtualisation Benefits**
  - Improved Resources Utilisation
  - Sharing Hardware Resources
  - Ease of Migration
  - Full VM Back-ups
  - Architecture Independence

© The Terma Group 2017

Use Cases (continued)

- **Virtualisation for Real-time Simulation**
  - Reproducibility (Archiving VMs)
  - Versioning of Tests (Versioning VMs/Snapshots)
  - Joint Execution of Multiple Simulations
    - Interconnect VMs on Physical Host
    - Reduce Network Impact
    - “Incompatible” Simulations (OS, Architecture, …)
  - Deployment and Sharing of Simulations
    - Ship VMs instead of Hardware

© The Terma Group 2017

Motivations Revisited

- **Independent Simulations**
  - Isolation
    - Motivation: Self-contained Setups, Host OS/CPU Architecture Independence, Reproducibility, Migration of Simulations
  - Parallelization
    - Motivation: Speed-up, Training, Improved Resource Utilization
- **Interdependent Simulations**
  - Multiple Systems or Components of Systems
    - Motivation: Integration
    - Varying Requirements regarding Synchronization (Systems vs. Components of a System)

© The Terma Group 2017

Problems?

- **“Realism” of Simulation**
  - Simulation of Critical Systems
  - Reliability of Results
  - Predictability
- **Impact of Virtualisation?**
  - Here: Real-time vs. Virtualisation
  - Influencing Factors?
  - Impact?
  - Assess “Reliability”/”Predictability”?

© The Terma Group 2017

Literature Review Overview (Virt. vs. RT)

- **Scheduler**
  - M. Lee, A. S. Krishnakumar, P. Krishnan, N. Singh, and S. Yajnik, “XenTune: Detecting xen scheduling bottlenecks for media applications,” in IEEE Global Telecommunications Conference (GLOBECOM 2010).
  - R. Ma, J. Li, L. Lin, and H. Guan, “DaSS: Dynamic time slice scheduler for virtual machine monitor,” in Algorithms and Architectures for Parallel Processing, Springer International Publishing, 2015.
  - …
- **IO, Memory, and IRQs**
  - C. Herber, A. Richter, T. Wild, and A. Herkersdorf, “Deadline-aware interrupt coalescing in controller area network (CAN),” in IEEE Intl Conf. on High Performance Computing and Communications, 2014.
  - …
- **CPU Pinning and Cache**
  - N. Mahmud, K. Sandstrm, and A. Vulgarakis, “Evaluating industrial applicability of virtualization on a distributed multicore platform,” in IEEE Emerging Technology and Factory Automation (ETFA), 2014.
  - …
- **Clock and Timekeeping**
  - B. Adamczyk and A. Chydzinski, “Achieving high resolution timer events in virtualized environment,” PLOS ONE, vol. 10, no. 7, 2015.
  - …

© The Terma Group 2017

Virt. vs. RT - Influencing Factors

- **Virtualisation Implementation**
  - KVM, Xen, VMware ESXi, …
- **Number and Placement of VMs**
  - Single, Same Core, Same Package, Different Package
- **Emulator/Simulator Implementation**
  - RTEMS, T-EMU, …
- **Stress Tests**
  - Dhrystone, Message Queue, Timer, …
- **Stressor Load**
  - CPU, Memory, HDD, Interrupt, …
- **Configuration**
  - Hardware (BIOS), Software
- **…**

© The Terma Group 2017

Test Suite: Motivation

- **Problem**
  - Number of Permutations of Influencing Factors
    - Executed about 1500 Tests during the Study
- **Solution**
  - Automation
    - Test Case Execution
      - Stress Test, Stressor, Measurements
    - Data Collection & Archival
    - Data Processing
    - Report Generation

© The Terma Group 2017

Test Suite: High-level Interaction Overview

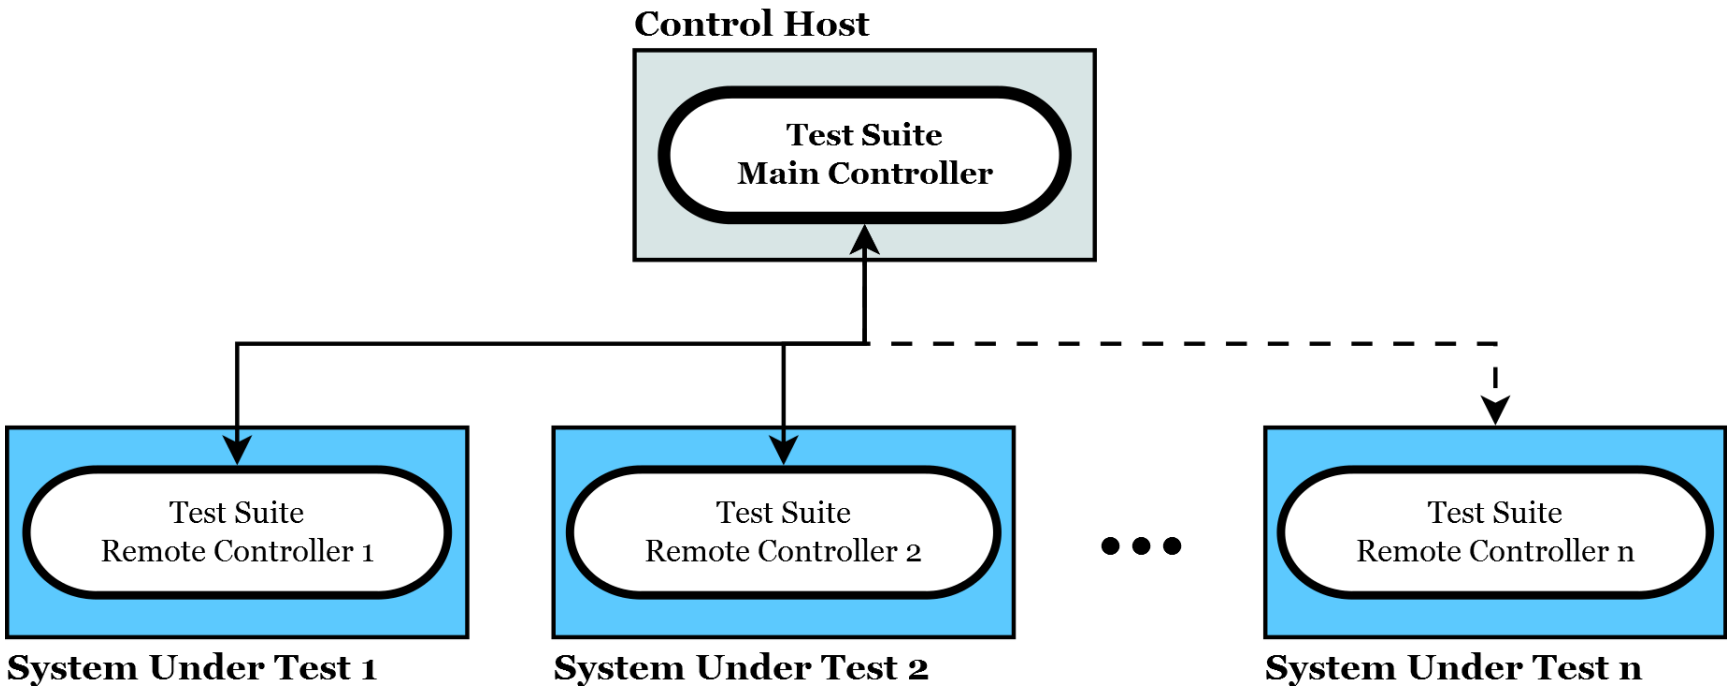

© The Terma Group 2017

Test Suite: More Detailed Overview

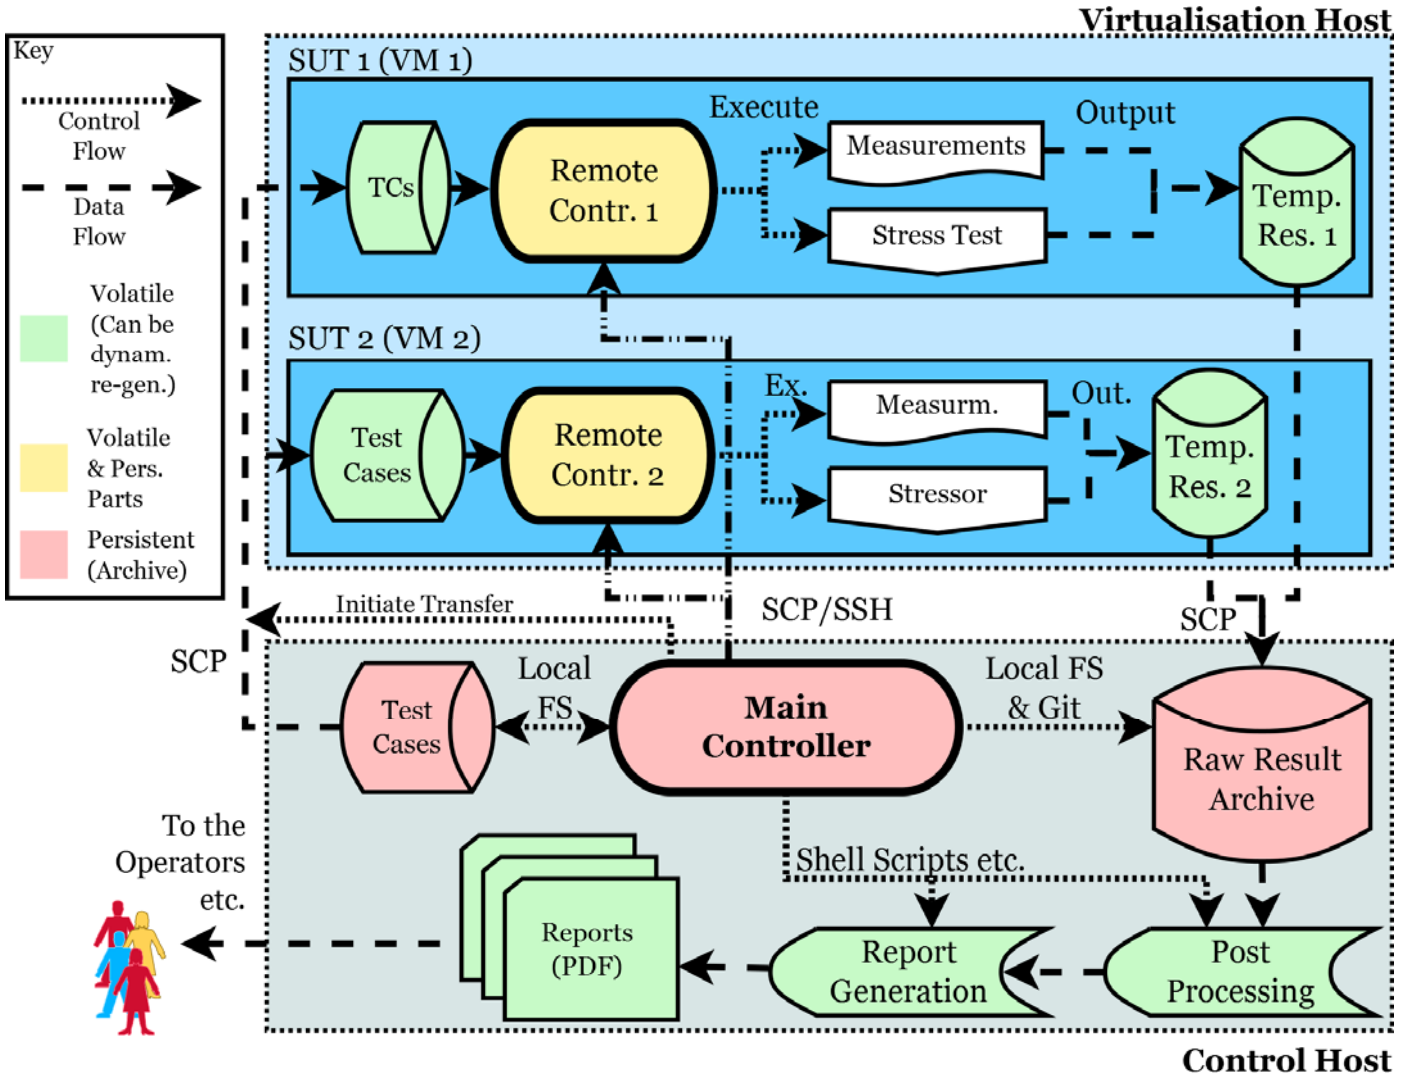

- **Automation, TC Execution, Rep. Gen., Modularity, Virt. Impl. Independent**

© The Terma Group 2017

Test Suite: Network Interconnection

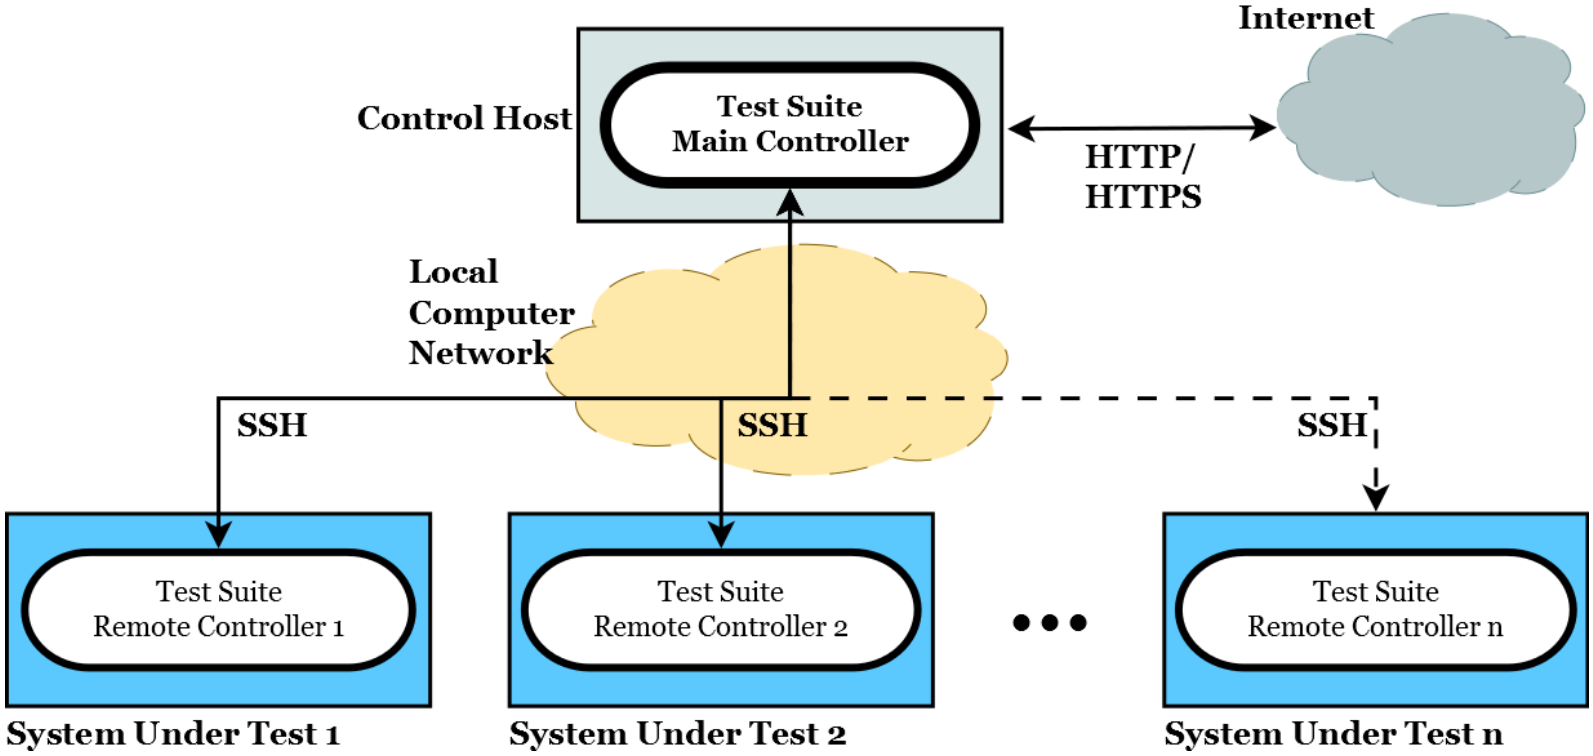

- Internet Connection: Only required for Control Host Installation
- Operation & SUT Base Installation: Only require SSH/SCP.

© The Terma Group 2017

Test Suite: Modularity & Extendibility

- **Measurement & Stress Tests**
  - Modular Approach
  - Dynamically Loaded
  - One Measurement/Stress Test per “Module”
  - Abstractions & Convenience Functionality
    - Execution, Post Processing, and Output Generation
- **Test Cases**
  - Use Pre-defined Measurements & Stress Tests
  - Configurable and Extendible
- **Virtualisation Implementation Independence**
  - Unified VM HDD Image for KVM, Xen, VMware ESXi
  - Recommendation: libvirt for Unified Tooling

© The Terma Group 2017

Test Suite: Measurement Example

```
# measurement_vmstat.py
class measurement_vmstat(MeasurementBase):
  def getInvocationCommandLine(self, destDir):
    return 'vmstat -n 1 > ' + self. getDst(destDir)

  def processRawData(self, sourceDir, tmpDir, destDir):
    tmpOutFile = self.getDst(tmpDir)
    dp.toUnifiedCsv(self.getDst(sourceDir), tmpOutFile)
    outSelector = {'cpu': ['us', 'sy', 'id', 'wa', 'st'],
                   'block_int_ctx': ['bi', 'bo', 'in', 'cs']}
    for outSel in outputSelector:
      dat=self.getDst(tmpDir, suffix = '_' + outSel + '.dat')
      dp.csvToGnuplot(tmpOutFile, dat, outputSelector[outSel])
      dp.renderGnuPlot(dat, self.getDst(destDir, '_' + outSel
+ '_line'))
```

© The Terma Group 2017

Test Suite: Test Case Configuration Example

```
# test_cases/example_tc.json
{"TestCaseName": "example_tc",
 "TestCaseDescription": "Simple TC Example.",
 "Hosts" : [{"HostAddress": "10.1.17.118",
             "User": "mcore",
             "Measurements": ["top", "vmstat"],
             "StressTests": ["dhrystone"],
             "TestDelayStart": 1,
             "TestDelayPre": 5,
             "TestDuration": 30,
             "TestDelayPost": 5,}]}
```

© The Terma Group 2017

Test Suite: Tests

- **T-EMU**
  - T-EMU configured for a LEON2 based system
  - Using ROM images built with rtems-4.8 and mkprom2
  - T-EMU plugins:
    - timereportplugin: Cyclic timereporting triggered by SRT.
    - timereportdevice: Memory mapped device model, accessed using timereporting.c (using sockets on host compiled RTEMS and the MMIO device on the emulator).
    - tmtc-link: Memory mapped and interrupt driven TMTC I/O model, sends and receives data to ground-sim.

© The Terma Group 2017

Test Suite: Tests (continued)

- **Stress Tests (RTEMS/T-EMU)**
  - dhrystone: Reports time every time the Dhrystone loop finishes.
  - rtems-irqlatency: Only emulator (host is handled by cyclictest)
  - rtems-mq: reports when data is received on RTEMS MQ.
  - rtems-timer: reports when timer triggers
  - …
- **Host only**
  - Stress Test
    - cyclictest
  - Stressor
    - stress-ng
      - CPU, RAM, HDD, Interrupt

© The Terma Group 2017

Test Suite: Time Reporting Server (TRS)

- **Ensures independent time reporting**
  - WCT on system under test
  - SRT on emulator
  - WCT on time reporting server host
- **UDP based protocol**
  - SUT reports to TRS.
  - TRS records reception time.
- **We are not looking at the exact times.**
- **Host compiled**
  - Uses direct sockets to send data.
- **Emulator**
  - Uses device model to sample time stamps at certain points.

© The Terma Group 2017

Test Suite: Report Generation

- **“Unified” Intermediate Data Format**
  - CSV Table
    - Can be easily used in, e.g.: R, Python, Excel, …
- **Transformation to Gnuplot “.dat” from Unified Format**
  - dp.csvToGnuplot(…)
  - Possibility to Select “Columns”
- **Pre-defined, Flexible Gnuplot Plots (“*.gpl” Files)**
  - Linechart, Boxplot, and Jitterplot with Linear & Y-Axis Log Scale
  - Support Arbitrary Number of Data Sets (Data “Columns”)
  - Default: PDF & PNG Output
- **Automatic Report Summary Generation**
  - One Large PDF-File
  - Based on LaTeX (pdflatex)

© The Terma Group 2017

Test Suite: Output Generation Example

```
dp.toUnifiedCsv(self.getDst(sourceDir), tmpOutFile)
outSelector = {'cpu': ['us', 'sy', 'id', 'wa', 'st'],
                    'block_int_ctx': ['bi', 'bo', 'in', 'cs']}
  for outSel in outputSelector:
    dat=self.getDst(tmpDir, suffix = '_' + outSel + '.dat')
    dp.csvToGnuplot(tmpOutFile, dat, outputSelector[outSel])
    dp.renderGnuPlot(dat, self.getDst(destDir, '_' + outSel +
'_line'))
    dp.renderGnuPlot(dat, self.getDst(destDir, '_' +
str(outSel) + '_line_log'), log = True)
    dp.renderGnuPlot(dat, self.getDst(destDir, '_' +
str(outSel) + '_box'), plotType='box')
    dp.renderGnuPlot(dat, self.getDst(destDir, '_' +
str(outSel) + '_box_log'), plotType='box', log = True)
    …
```

© The Terma Group 2017

Test Suite: Output Plots Example

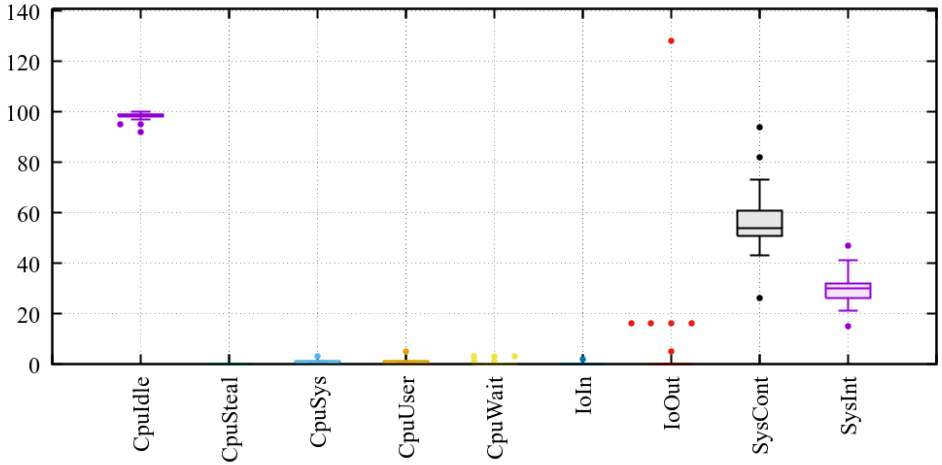

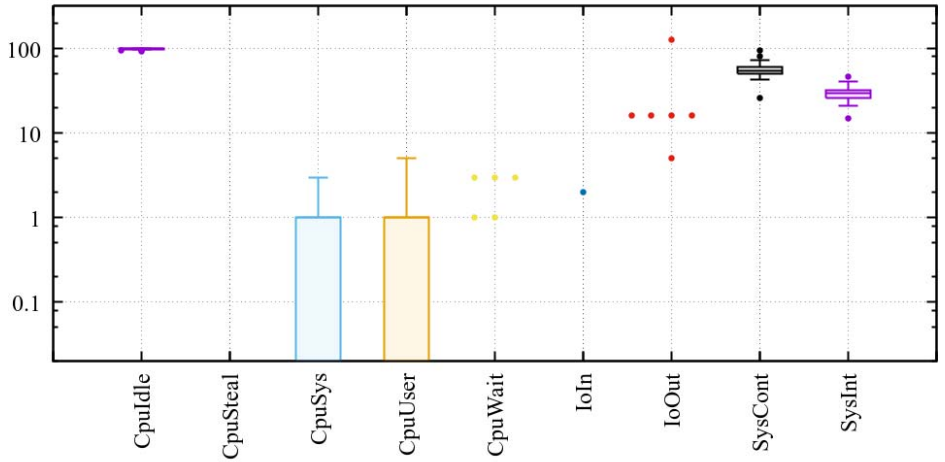

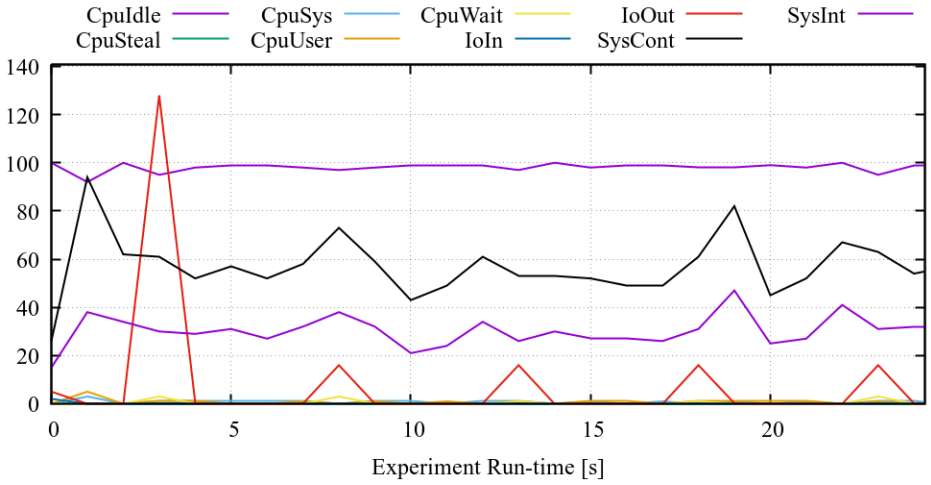

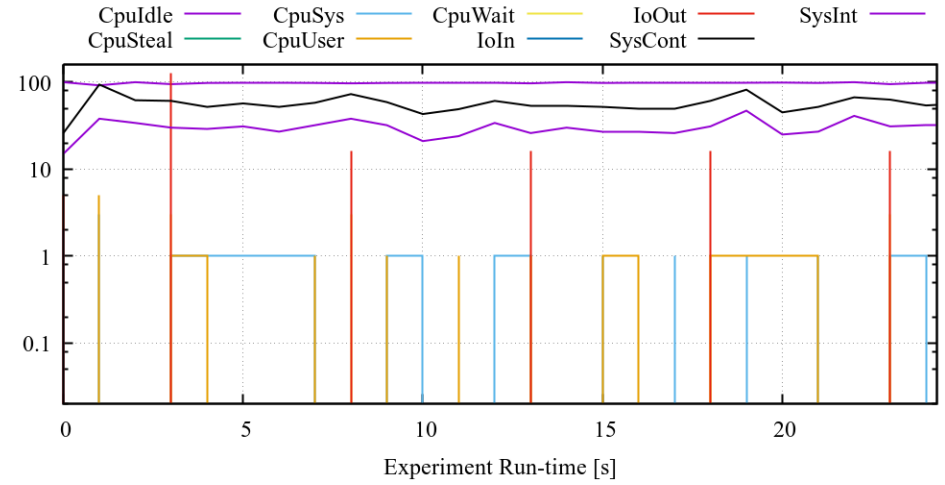

© The Terma Group 2017

Results: Preface

- **Timekeeping as Example**
  - Two Timers
  - Started at “same” Time
  - Time Difference between Triggering of Timers
- **Two Timestamps as Example**
  - Reported Simulated Real-time (**repSrtT**)
  - Report Reception Time (**rcvT**)

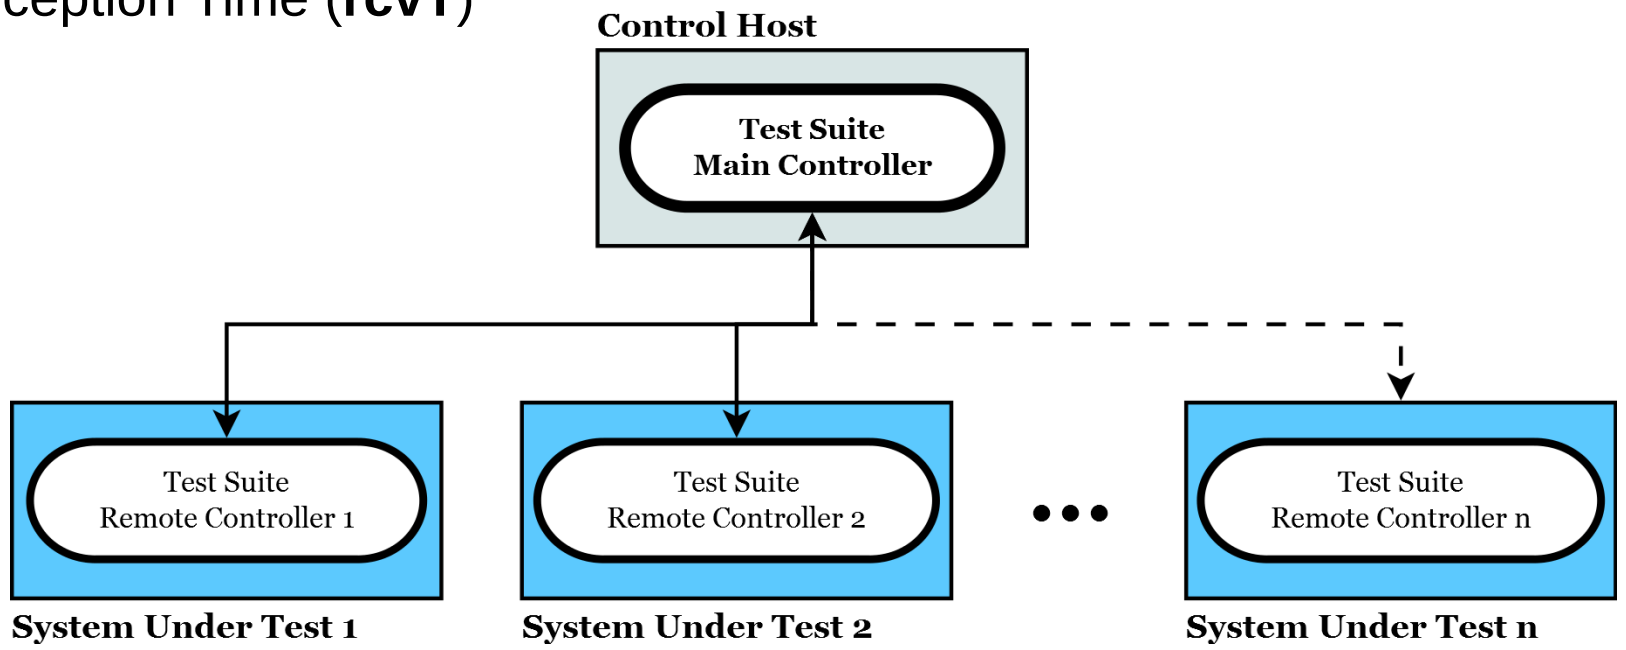

© The Terma Group 2017

Results, Time Keeping, Single Host

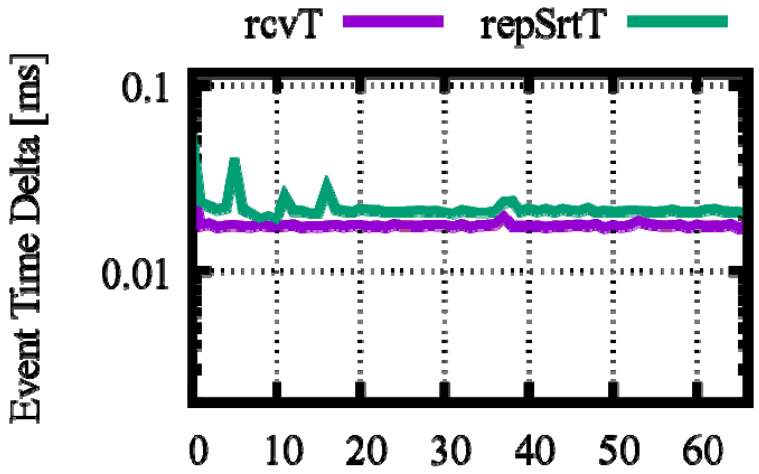

Bare metal

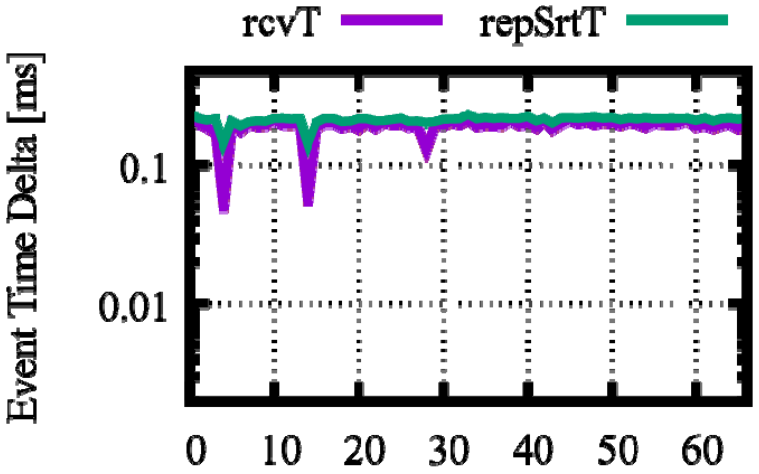

KVM

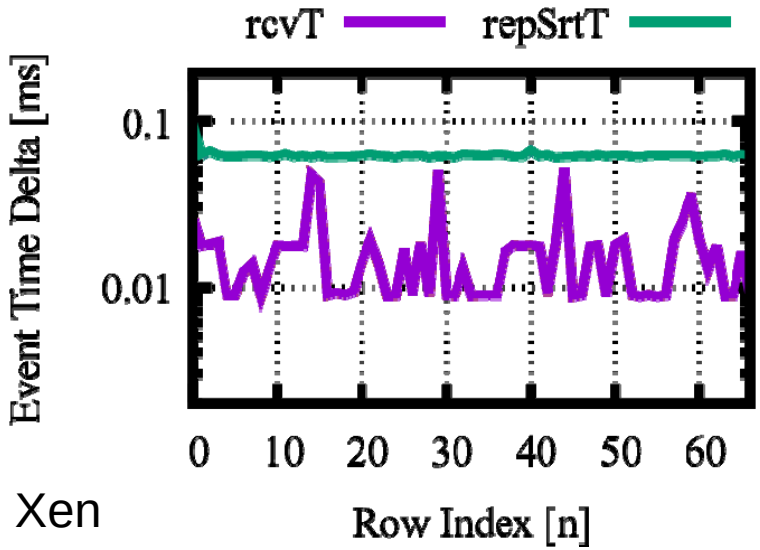

Xen

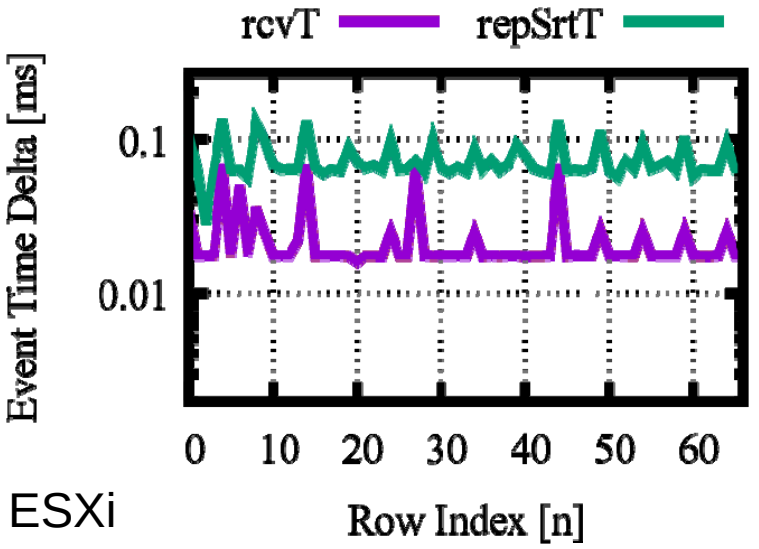

ESXi

© The Terma Group 2017

Results, Time Keeping, XEN Dual Host

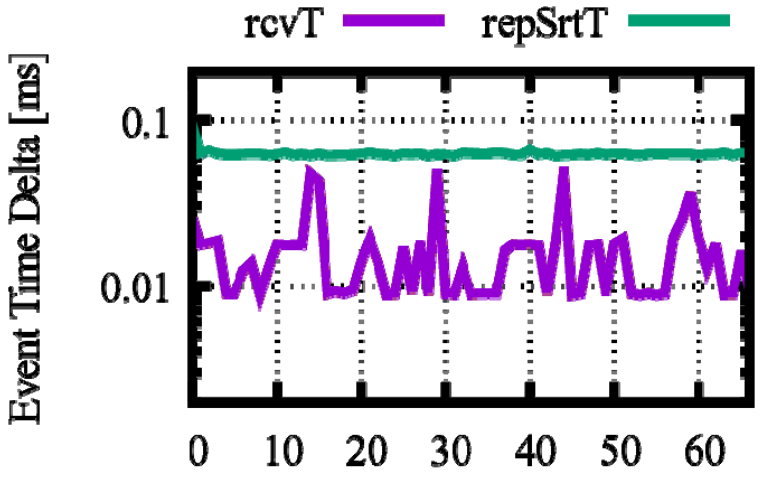

Single host

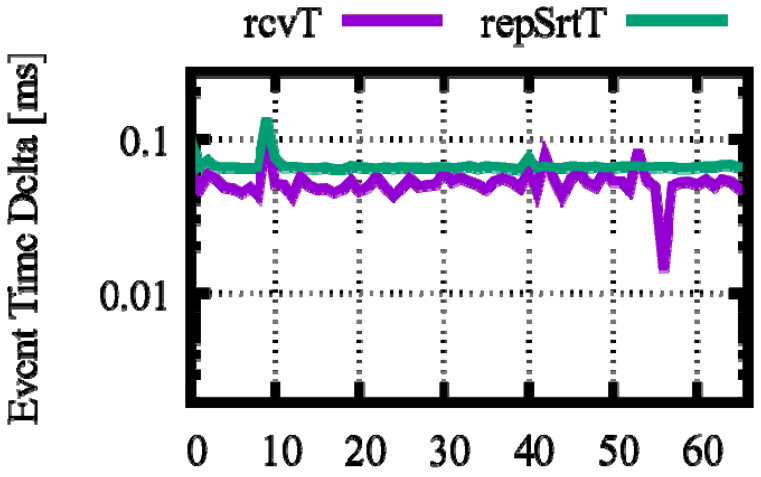

Different Package

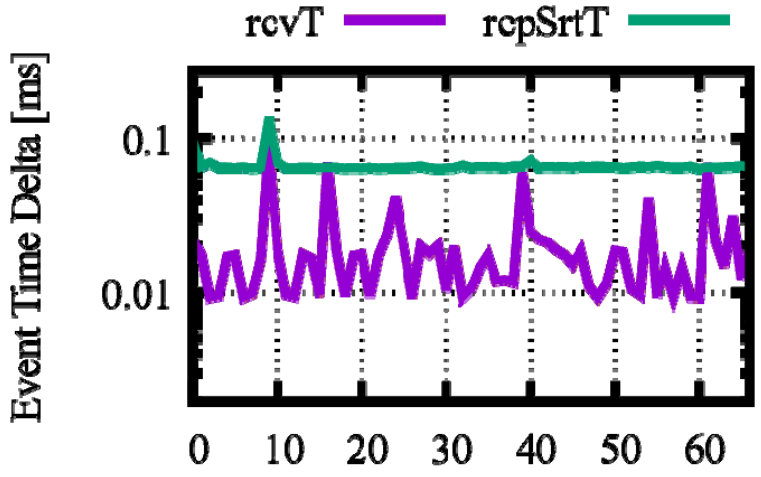

Same Package

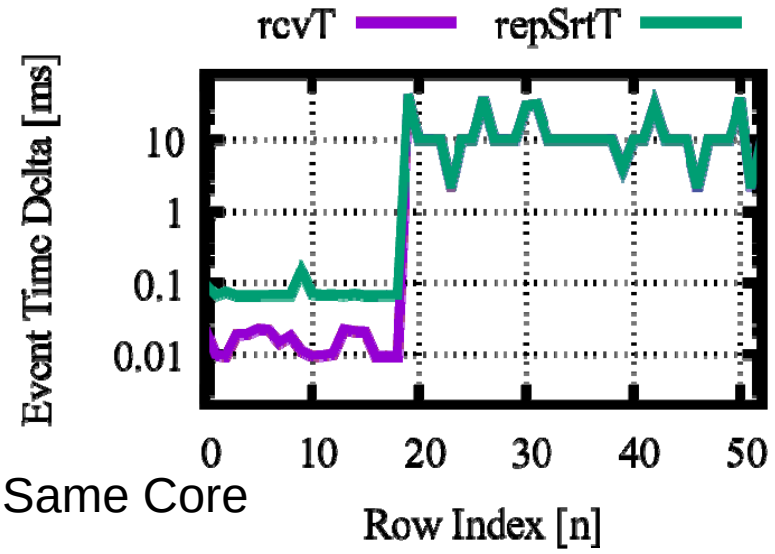

Same Core

© The Terma Group 2017

Results, Time Keeping, Dual H., Same Core

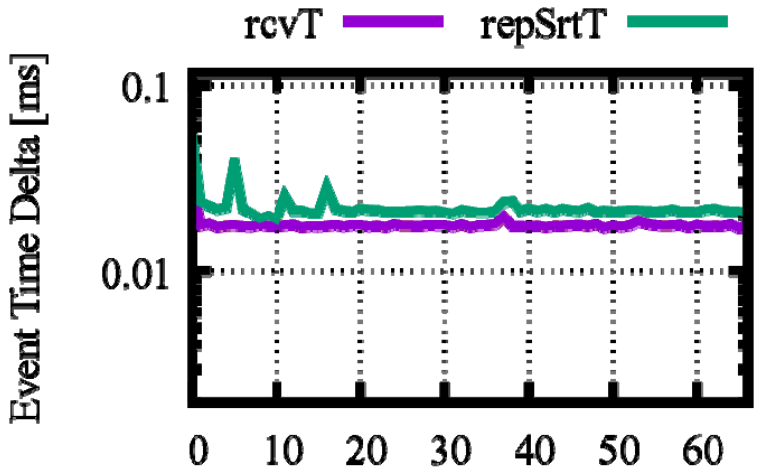

Sing. H., Bare Metal

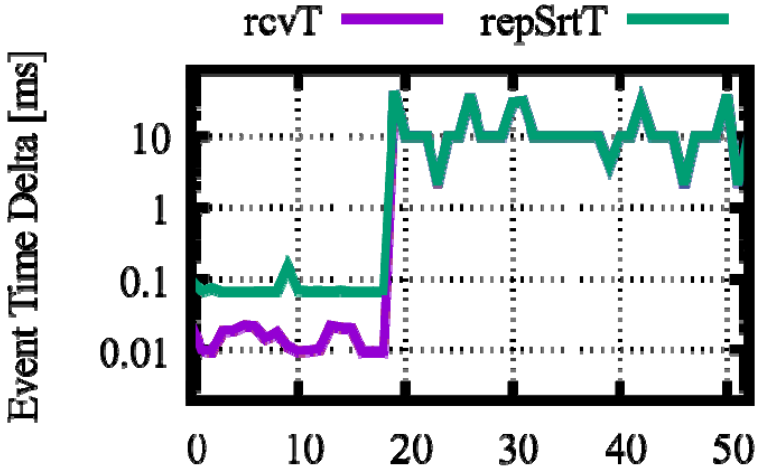

Xen

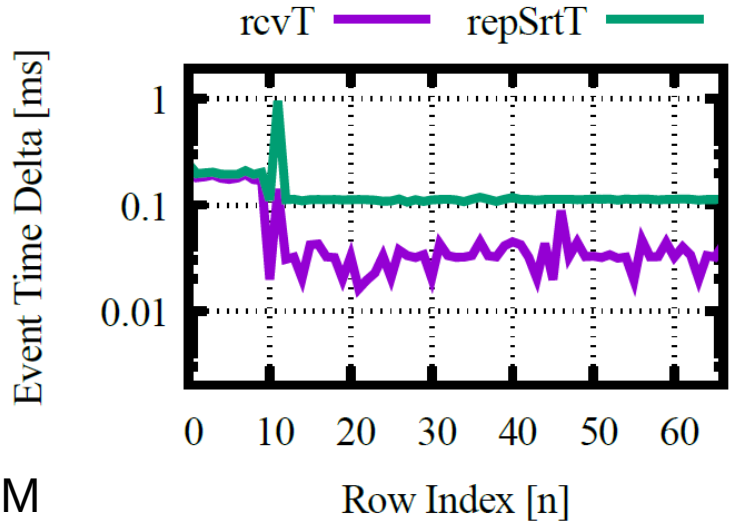

KVM

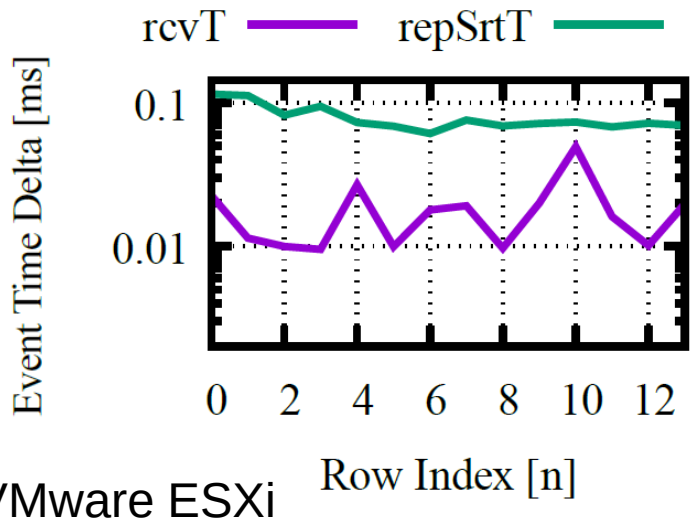

VMware ESXi

© The Terma Group 2017

NUMA Effects

- **Shared Memory Bandwidth**
  - CPU(s) at same Memory Controller
- **Memory Access via „Remote“ Memory Controller**
  - CPUs at different Memory Controller
  - Access Memory from other CPU
- **Impact depends on scenario.**
  - System Level, Component Level, CPU Core Level
  - Memory Throughput / Utilization
  - …

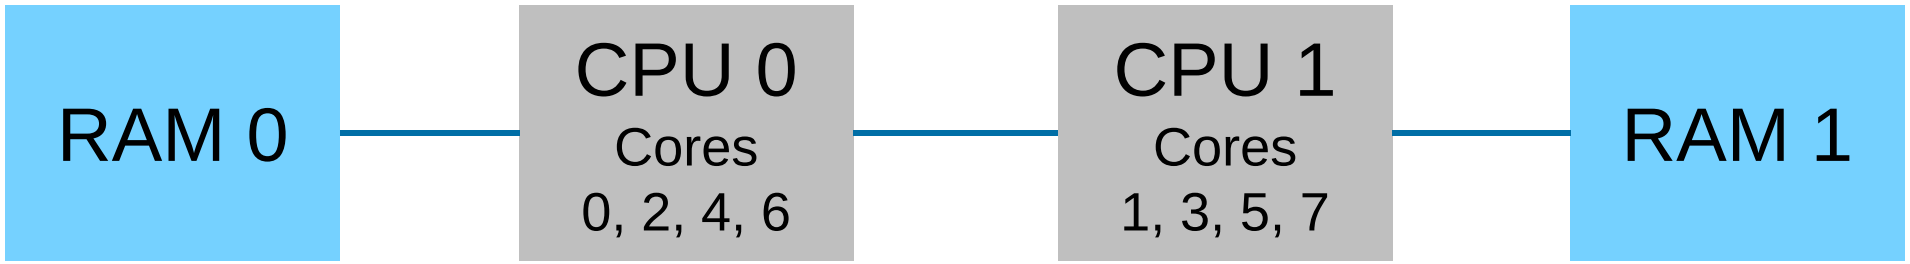

© The Terma Group 2017

Guidelines

- **Hardware**
  - Supports virtualisation technology, such as AMD-V or Intel-VT
  - Prefer more CPU sockets over cores per CPU
  - Avoid oversharing among VMs (RAM, CPU, HDD, …).
- **VM Configuration**
  - CPU Pinning
  - Avoid Virtual CPUs (Hyperthreading etc.)
  - Be aware of NUMA.
  - I/Os: pass-through when possible
  - Isolate VMs as far as possible.
- **Software**
  - Network Time Synchronisation (NTP, PTP)
  - Real-time Kernel and Software (Host & Guest)

© The Terma Group 2017

Further Results

- **Libvirt**
  - Virtualisation Implementation Independence
  - Unified Interface & Configuration File Format
    - Here: KVM, Xen, VMware ESXi
  - Scriptable -> Automation
- **Unified VM Hard Disk Image**
  - KVM & Xen same File
    - Qcow2 Format
  - VMware ESXi Conversion
    - Qcow2 to VMDK Format
  - Share & Migrate VMs

© The Terma Group 2017

Further Results (continued)

- **Record as much as possible.**
  - As long as recoding does not influence the results.
- **Git**
  - Raw Results
    - Automation
  - Output Data
  - Test Suite & Test Cases
- **Report Generation**
  - Automation
  - Better Overview of Results

© The Terma Group 2017

Summary & Conclusion

- **Virtualisation for RT Simulation/Emulation**
  - Benefits
  - Potential Problems
  - Large Number of Influencing Factors
  - Impact Depends on Use Case
- **Test Suite**
  - Help to Assess Concrete Suitability for RT Sim./Emul.
  - High-degree of Automation
    - Execution, Data Collection, Report Generation
  - Can also be used in „Production“
    - Assess „Quality“ of Setup
    - Monitor Running Experiments
- **Results & Guidelines**

© The Terma Group 2017

End

Thank you very much for your attention.

Questions?

© The Terma Group 2017

Meet us at

www.terma.com

www.terma.com/press/newsletter

www.linkedin.com/company/terma-a-s

www.twitter.com/terma_global

www.youtube.com/user/TermaTV