OCTOBER 2023

Virtualization best practices and troubleshooting

Bob Hunter, Principal GCS Technical Account Manager

David Gardner, SR GCS Technical Account Manager

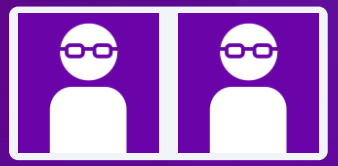

Bob Hunter

Principal GCS Technical Account Manager

AVEVA

bob.hunter@aveva.com

David Gardner

SR GCS Technical Account Manager

AVEVA

david.gardner@aveva.com

Agenda

- Summary or what are we doing here?
- Justification or why virtualize?
- Machine roles
- Recommendations
- Issues

What are we doing here?

Virtualization practices

1. Why use virtual computers
2. What are the machines that can / should be virtualized
3. What are some recommendations / best practices
4. Issues seen when using virtualization in Operation

Why virtualize

Why Is Virtualization Important?

- Consolidate servers
- Simplify management
- Minimize downtime
- Corporate/IT initiatives (SaaS/Cloud)

Machine roles

Machine roles and resource estimates

“Medium sized” system

Role	CPU Cores	RAM (GB)	HDD (GB)
GR Server	8-16	8-12	200-500
Historian Server	8-16	8-12	250-500
Engineering Workstation	8	4	200
Visualization – single user	8	4	200
Visualization – RDP Client	4	2	N/A
AOS / DAS	8-16	8-12	200-500

**refer to Readme.txt on the installation media*

AVEVA

Design – OS Templates

- Build your own library from scratch
- Protect your templates with AV Software

Product	Client OS (x64 only)	Server OS (x64 only)
AVEVA System Platform (Application Server)	• Window 10 or later • LTSC Version 1809 or later • Enterprise and IiT Enterprise • SAC Version 20H2	• Windows Server 2016 or later • Standard, Data Center, and IoT • SAC versions or later Windows Server OS are not Supported
AVEVA InTouch HMI		
AVEVA Historian		

**refer to Readme.txt on the installation media*

AVEVA

Key recommendations

Resource recommendations

Key recommendations to help ensure a well running VM system

Target CPU utilization	Deploy virtual machines to a host with the goal of keeping server processor utilization at 50% or less. Plan on not more than 2 App engines per vCore
Use memory reservations	Utilize CPU and memory reservations for the virtual machine running AVEVA software. Virtual machines with reservations get preferential treatment by the scheduler in the host.
High-speed processing	To ensure proper function of Realtime critical applications do not over-commit resources on the host and keep processor affinity locked. For example, if a physical unit was going to be 2 cores and 2GB of RAM, dedicate this to the virtual machine as well and make it the minimum.
Host machines resources	Should always have at least 25% more resources than what the guest machines would utilize if they used 100% of their assigned resources

Issues

Issues using Virtualization

- IT versus OT goals
- Clock drift
- Snapshots
- Recovering from machine failure

Questions?

Please wait for the microphone.

State your name and company.

Please remember to…

Navigate to this session in the mobile app to complete the survey.

Thank you!

AVEVA

Additional Resources

- AVEVA™ System Platform 2023 media
  - Readme.html
  - SP_Virtual_Implementation.html

The following guidelines are provided for reference only. The system configuration required for your application will depend on multiple factors, including but not limited to the size and complexity of the application, and the features and components used.

Application	Level	Logical Processors [1]	RAM [3]	Free Disk Space [2, 3]	Network Speed
Application Server Nodes [5]					
Small Application (1 - 25,000 I/O per Node)	Minimum	4	2 GB	100 GB	100 Mbps
	Recommended	8	4 GB	200 GB	1 Gbps
Medium Application (25,000 - 50,000 I/O per Node)	Minimum	8	8 GB	200 GB	1 Gbps
	Recommended	16	12 GB	500 GB	1 Gbps
Large Application (> 50,000 I/O per Node)	Minimum	16	16 GB	500 GB	1 Gbps
	Recommended	32	24 GB	1 TB	Dual 1 Gbps
Galaxy Repository Nodes					
Small Galaxy (1 - 50,000 I/O per Node)	Minimum	4	2 GB	100 GB	100 Mbps
	Recommended	8	4 GB	200 GB	1 Gbps
Medium Galaxy (50,000 - 200,000 I/O per Node)	Minimum	8	8 GB	200 GB	1 Gbps
	Recommended	16	12 GB	500 GB	1 Gbps
Large Galaxy (> 200,000 I/O per Node)	Minimum	16	16 GB	500 GB	1 Gbps
	Recommended	32	24 GB	1 TB	Dual 1 Gbps
Historian Server Nodes					
Small Historian (1 - 50,000 Historized Tags per Node)	Minimum	4	2 GB	100 GB	100 Mbps
	Recommended	8	4 GB	200 GB	1 Gbps
Medium Historian (50,000 - 200,000 Historized Tags per Node)	Minimum	8	8 GB	200 GB	1 Gbps
	Recommended	16	12 GB	500 GB	1 Gbps
Large Historian (> 200,000 Historized Tags per Node)	Minimum	16	16 GB	500 GB	1 Gbps
	Recommended	32	24 GB	1 TB	Dual 1 Gbps
Thin Client Nodes					
RDP clients, InTouch Access Anywhere, mobile devices	Minimum	2	512 MB	N/A	100 Mbps
	Recommended	4	2 GB	N/A	100 Mbps
Client Nodes					
WindowViewer, OMI ViewApp, Historian Client, Remote IDE	Minimum	4	1 GB	100 GB	100 Mbps
	Recommended	8	4 GB	200 GB	1 Gbps
Remote-Desktop Server Nodes					
Basic RDS, InTouch Access Anywhere Server Supports up to 15 concurrent remote sessions	Minimum	8	8 GB	200 GB	1 Gbps
	Recommended	16	12 GB	500 GB	1 Gbps
Large RDS, InTouch Access Anywhere Server Supports up to 30 concurrent remote sessions	Minimum	16	16 GB	500 GB	1 Gbps
	Recommended	32	24 GB	1 TB	Dual 1 Gbps
All-In-One Nodes [4]					
Small Application - All products on one node: 1,000 I/O max	Minimum	8	8 GB	200 GB	100 Mbps
	Recommended	12	12 GB	500 GB	1 Gbps
Medium Application - All products on one node: 20,000 I/O max	Minimum	12	16 GB	500 GB	1 Gbps
	Recommended	16	32 GB	1 TB	1 Gbps
Large Application [6] - All products on one node: 100,000 I/O max	Minimum	20	32 GB	2 TB	1 Gbps
	Recommended	24	64 GB	4 TB	1 Gbps

1) To calculate the number of logical processors: multiply the number of physical cores by the number of threads each core can run. A four core CPU that runs two threads per core provides eight logical processors. The terms "Hyper-Threading and "simultaneous multithreading" (SMT) are also used to describe logical processors.
2) SSD drives are highly recommended.
3) In redundant environments, increase CPU and RAM to maintain a maximum of 40% typical resource utilization.
4) For optimal performance of all-in-one nodes, a high clock speed [illegible] is recommended.

5) For Application Server platform nodes, it is recommended that you deploy no more than two AppEngines per logical processor (typically one primary AppEngine and one backup).
6) For large applications on all-in-one nodes, dual XEON processors are recommended.

Supported Operating Systems at Time of Release

The latest product information for each System Platform product is listed in the AVEVA Global Customer Support GCS Technology Matrix. Each link includes:

- General information about the selected product, such as the version number and release date
- Operating system requirements
- Microsoft SQL Server requirements
- Virtualization software compatibility
- Information about interoperability (which AVEVA products a specific product works with)
- Information about coexistence (which AVEVA products can be installed on the same node)

Note: Only Windows Server versions licensed under Microsoft's long term servicing channel (LTSC) are supported. Versions of Windows Server licensed under the Semi-Annual Channel (SAC) do not include Desktop Experience, and therefore, are not supported. For example SAC versions 1709, 1903, and 2003 are not supported.

64-bit only

System Platform Product / Component	SAC Windows 10/11 Pro/ Enterprise/IoT Enterprise	LTSC Windows 10/11 Enterprise/IoT Enterprise	LTSC Windows Server See Windows Server Notes
Application Server 2023 • Galaxy Repository • AppEngine / Platform / Bootstrap • IDE (including remote) • AVEVA OMI ViewApp (run time)	Windows 10 20H2 Windows 10 21H1 Windows 10 21H2 Windows 11 21H2 Windows 11 22H2 Includes Professional, Enterprise and IoT Enterprise.	Windows 10 1809 Windows 10 21H2 Windows 11 Includes Enterprise and IoT Enterprise	2016 Standard and Data Center 2016 IoT 2019 Data Center, Essentials, and Standard 2019 IoT 2022 Standard and Data Center
InTouch HMI 2023 • WindowMaker (no Modern apps) • WindowMaker (Modern apps) • WindowViewer (run time) / InTouchViewApp • InTouch Web Client (run time)			
InTouch for System Platform 2023 • WindowMaker (Managed Apps) • WindowViewer (run time) / InTouchViewApp			
Historian Server 2023			
InTouch Access Anywhere Server 2023 • ITAA Client (HTML5 browser) InTouch Access Anywhere Secure Gateway 2023			
Other System Platform 2023 components / products • Historian Client 2023 • OI Gateway • System Management Server • AVEVA Licensing • System Monitor Agent / Manager			

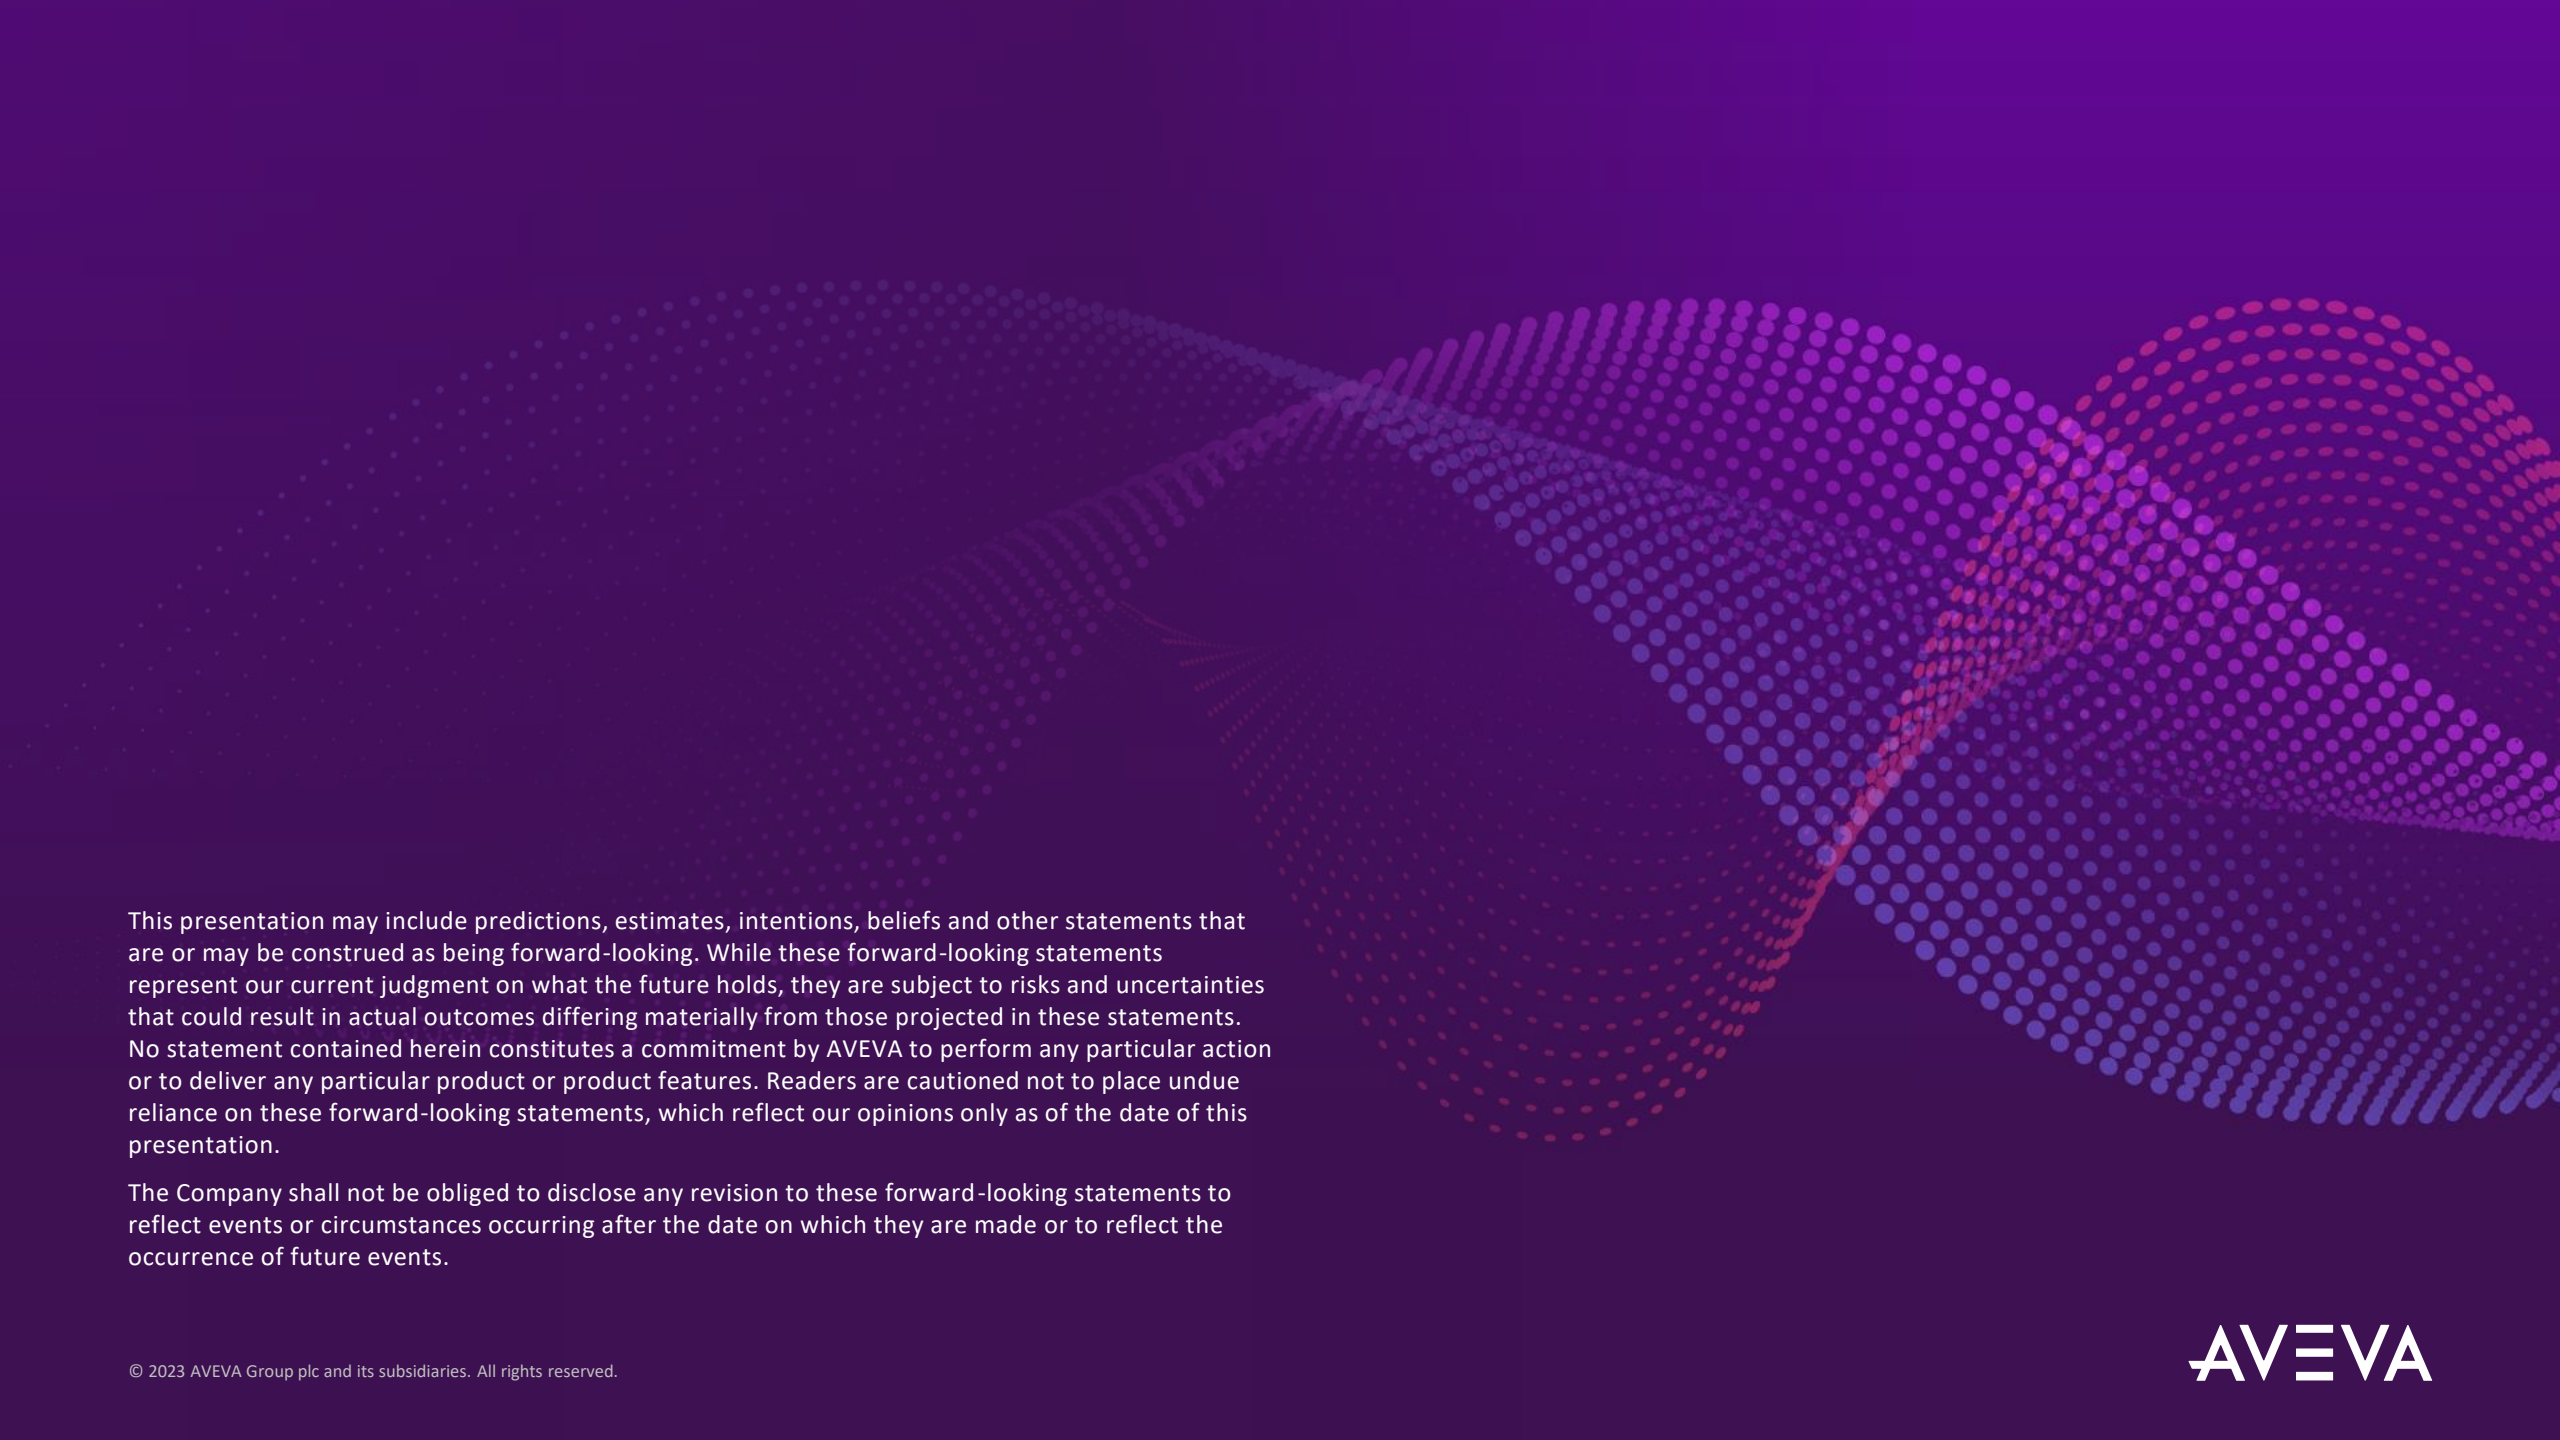

This presentation may include predictions, estimates, intentions, beliefs and other statements that are or may be construed as being forward-looking. While these forward-looking statements represent our current judgment on what the future holds, they are subject to risks and uncertainties that could result in actual outcomes differing materially from those projected in these statements. No statement contained herein constitutes a commitment by AVEVA to perform any particular action or to deliver any particular product or product features. Readers are cautioned not to place undue reliance on these forward-looking statements, which reflect our opinions only as of the date of this presentation.

The Company shall not be obliged to disclose any revision to these forward-looking statements to reflect events or circumstances occurring after the date on which they are made or to reflect the occurrence of future events.

linkedin.com/company/aveva

@avevagroup

ABOUT AVEVA

AVEVA is a world leader in industrial software, providing engineering and operational solutions across multiple industries, including oil and gas, chemical, pharmaceutical, power and utilities, marine, renewables, and food and beverage. Our agnostic and open architecture helps organizations design, build, operate, maintain and optimize the complete lifecycle of complex industrial assets, from production plants and offshore platforms to manufactured consumer goods.

Over 20,000 enterprises in over 100 countries rely on AVEVA to help them deliver life's essentials: safe and reliable energy, food, medicines, infrastructure and more. By connecting people with trusted information and AI-enriched insights, AVEVA enables teams to engineer efficiently and optimize operations, driving growth and sustainability.

Named as one of the world's most innovative companies, AVEVA supports customers with open solutions and the expertise of more than 6,400 employees, 5,000 partners and 5,700 certified developers. The company is headquartered in Cambridge, UK.

Learn more at www.aveva.com

AVEVA