

VGOS Data Recorder Comparison

Compiled by Bill Petrachenko (NRC Canada) & Alan Whitney (MIT Haystack Observatory)

June 2013

Introduction.

The minimum VGOS data rate requirements are 16 Gbps 'burst' for 10 sec with ~4 Gbps sustained average rate. This is based on a typical VGOS schedule which alternates between on-source and slew/wait/cal periods with antennas being on source about 25% of the time (but could be greater under some circumstances) and the network slewing to the next source the rest of the time. The maximum on-source period for normal operations is set to 10 s but could be longer at times to achieve greater sensitivity.

At the present time only two VLBI data recorders fully satisfy VGOS operational requirements, the Mark 6 (already available) and the K6/iDAS (still under development). Because the K6/iDAS does not support removable media, it can only be used for e-transfer and not for physical transport of media. See Table 1 for more details.

Another pair of recorders, the Mark 5C and K5/VSI do not themselves meet VGOS data rate requirements. They are however close enough that if preceded by an external RAM buffer or combined in a group of four (plus some format conversions for the K5/VSI), they would become compliant. This solution is neither cost effective nor convenient and hence the Mark 5C and K5/VSI are not recommended for use in VGOS. See Table 2 for more details.

For completeness, other VLBI data recorders that are still in use but not suitable for VGOS are summarized in Tables 3a and 3b.

Fully Compliant VGOS Data Recorders, Table 1

Mark 6. The Mark 6 recorder is an outgrowth of the Haystack/Conduant Mark 5 development. Like its predecessor, the Mark 5C, I/O is based on fixed length Ethernet packets (e.g. VDIF), with the output being an exact replica of the input. Like all Mark 5/6 systems, recording media is removable and packaged as 8-disk SATA modules. Mechanically the Mark 6 includes a system enclosure, a disk module enclosure, and a cable management panel, all produced by Conduant. In terms of data electronics there has been a marked shift toward COTS components with only the disk-module-related hardware produced by Conduant; the motherboard (MB), NIC, and disk interface cards are all COTS. In contrast to the Mark 5 family that uses a proprietary Conduant media file-system format, the Mark 6 uses standard Linux files with 'scattered' data. All software has been developed at Haystack and is open source. The Mark 6 is ideally suited for use in VGOS. It has four 10GigE inputs (one per VGOS band), a maximum aggregate input data rate of a minimum of 16 Gbps into ~64GB of high-speed RAM, and a sustained record rate per installed disk module of ~7 Gbps (for mechanical disks; much higher for SSDs). Hence, it is possible to achieve a sustained record rate significantly greater than the VGOS requirement (~4 Gbps) using only a single 8-disk module, thus minimizing shipping costs. In order to support sustained rates up to 16Gbps or to support unattended operation over longer periods at lower rates, as many as four 8-disk media modules can be simultaneously mounted. Options are available to upgrade existing Mark 5 systems to Mark 6, as well as to upgrade existing Mark 5 SATA disk modules to Mark 6 compatibility.

K6/iDAS. The K6/iDAS means the "K6 intelligent Data Acquisition and Streaming" system, which is the name of the sampler with DBBC for VGOS operation developed in Japan. There is no dedicated recorder to the K6/iDAS sampler because it allows us to use itself with a fully COTS server, 10GigE-network switch, and storage system as recorder. You can adopt a server with a 10GigE network interface and a large amount of RAID as a recorder for K6/iDAS. GSI adopted IBM System x3650 M3, Arista 7050S64-T, Data Direct Networks (DDN) Storage with StorNext to be compliant with a sustained record rate of 32 Gbps, which leaves room for future enhancement of the VGOS requirements. The recording system adopted by GSI does not support removable media so it can only be used for e-transfer and not for physical transport of media. It also does not support the IVS standard VSI-S control protocol but is controlled directly by logging into the system.

| | Mark 6 | K6/iDAS (the recording system adopted by GSI) |
|---|--|--|
| VGOS Compatibility | Full | eVLBI (yes) Media physical transport (no) |
| Data Input Connection | 4x10GigE | 2x10GBASE-SR |
| Max data-rate per input connection | 8 Gbps | 8 Gbps |
| Max aggregate input data rate | 16 Gbps | 32 Gbps (4 sampler units/8 servers) |
| Max sustained record rate | ~7 Gbps x # of media packages installed (for hard disks); much faster for SSDs | 32 Gbps |
| RAM buffer size | 64 Gbytes (~30 s at 16 Gbps) | N. A. |
| Removable media? | Y | N |
| Recording media packaging | Mk6 8-disk SATA module | SATA & SAS Linux RAID with StoreNext |
| Max # of media packages per system | 4 | 1 |
| # media packages for max recording data rate | 2 to 4 depending on speed of disks | 1 |
| Input data format | VDIF or Mk5B Ethernet packets | VDIF (UDP/IP) |
| Record-data format | VDIF | VDIF |
| Media file-system format | Linux files w/'scattered' data | Linux files |
| Resilient to slow recording media? | Y | Y |
| Max aggregate playback rate | Up to 16 Gbps (depends on playback system details) | 32 Gbps |
| Playback data interface type | 10GigE or Infiniband | 1x10GigE x 16 servers |
| Playback data format | Recorded data payload in Ethernet frame | VDIF |
| OS type | Linux | Linux |
| Software developer ('proprietary' or 'open') | Haystack (open) | GSI/Nitsuki (open by GSI) |
| Command interface connection | GigE | GigE |
| Command language | VSI-S | Direct control by logging into the system |
| Current status | Available | Under development, partially available |
| Cost (USD) | ~\$14K USD | depend on rate and capacity |
| Contact | <mark6@haystack.mit.edu> | Shinobu Kurihara <skuri@gsi.go.jp> |

Table 1. VGOS Data Recorders – systems fully meeting the VGOS specification. However, the K6/iDAS does not support removable media and hence only works at stations dedicated to eVLBI.

Potentially Compliant (but not recommended) VGOS Data Recorders, Table 2

Mark 5C. The Mark 5C recorder is an outgrowth of the Haystack/Conduant Mark 5 development. In contrast to its predecessors, I/O is based on fixed length Ethernet packets (e.g. VDIF) with the output being an exact replica of the input. Like all Mark 5 systems, recording media is removable and packaged as Mk5 8-disk PATA or SATA modules. Mechanically the Mark 5C includes a system and disk-module enclosure produced by Conduant. Its data electronics include a standard COTS motherboard and disk array interface, 10GigE daughterboard and data distributor produced by Conduant. The Mark 5C could conceivably be adapted to support VGOS if four Mark 5's were used at each station or if it were preceded in the signal chain by a burst mode RAM buffer, but these options are neither cost effective nor convenient; hence the Mark 5C is not recommended for VGOS.

K5/VSI. The K5/VSI is a VSI-H based VLBI recorder using K5 data-frame format. It is housed in a standard PC enclosure with a proprietary PC-VSI board from Digital Link and a COTS NIC and disk interface. The K5/VSI could conceivably be adapted to support VGOS if four K5/VSI's were used at each station if it were preceded in the signal chain by a burst mode RAM buffer, but these options are neither cost effective nor convenient. Hence the K5/VSI is not sufficient for VGOS. The K5/VSI does not support removable media so it can only be used for e-transfer and not for physical transport of media. Finally, it does not support the standard IVS VSI-S control protocol but is controlled directly by logging into the system at present. Since the recording is totally controlled by the software on standard linux, this feature can be improved by software cording on the linux OS without any difficulty.

| | <i>Mark 5C</i> | <i>K5/VSI</i> |
|---|--|--|
| VGOS Compatibility | Requires 4 systems or the addition of a burst mode buffer. | Requires 4 systems or the addition of a burst mode buffer, plus VSI-H to VDIF converter. |
| Data Input Connection | 10GigE | VSI-H |
| Max data-rate per input connection | 4 Gbps | 2 Gbps |
| Max aggregate input data rate | 4 Gbps | 4 Gbps (with 2 PC's) |
| Max sustained record rate | 4 Gbps | 4 Gbps (with 2 PC's) |
| RAM buffer size | N.A. | N.A. |
| Removable media? | Y | N |
| Recording media packaging | Mk5 8-disk PATA or SATA module | RAID array in commercial pkg. |
| Max # of media packages per system | 2 | N.A. |
| # media packages for max recording data rate | 2 | 2 (with 2 PC's) |
| Input data format | Any fixed-length Ethernet package (e.g. VDIF) | VSI-H |
| Record-data format | Ethernet-frame data payload (typically Mk5B or VDIF) | VDIF, K5/VSSP, Mark5B |
| Media file-system format | Proprietary | Linux files |
| Resilient to slow recording media? | Y | Y |
| Max aggregate playback rate | ~1 Gbps | Depends on speed of PC ~6 Gbps |

| | | |
|---|-------------------------------------|---|
| Playback data interface type | GigE | 1GigE or 10GigE |
| Playback data format | Same as input format | 1GigE or 10GigE |
| OS type | Linux | Linux |
| Software developer ('proprietary' or 'open') | Conduant (prop) and Haystack (open) | Data Link (prop) NICT (open) |
| Command interface connection | GigE | GigE |
| Command language | VSI-S | Direct control by logging into the system |
| Current status | Available | Available |
| Cost (USD) | ~23K USD | ~20K USD |
| Contact | <mark5@haystack.mit.edu> | Mamoru Sekido <sekido@nict.go.jp> |

Table 2. Potentially compliant (but not recommended) VGOS Data Recorders.

VLBI Data Recorders that are not VGOS compliant, Table 3a

Mark 5A/B. The Mark 5A recorder was the first entry in the Mark 5 family and as such was designed to have minimal disruption from existing tape based systems. Hence it used a VLBA/Mark4 input data format and a recorded data-frame format that is based on the VLBA/Mark4 tape systems. The Mark5B represents a move towards IVS standards and uses a VSI-H input data format and a Mark 5B data-frame format. In both cases, recording media is removable and packaged as Mk5 8-disk PATA or SATA modules and the media file-system format is proprietary to Conduant. Mechanically, the Mark 5A/B uses a system and disk module enclosure produced by Conduant. Its data electronics include a standard COTS motherboard with proprietary disk array interface and data distributor produced by Conduant. Mark 5A/B data rates are too low to be useful for VGOS and the systems are no longer available.

| | Mark 5A | Mark 5A+ | Mark 5B | Mark 5B+ |
|---|-----------------------------------|-----------------|----------------|-----------------|
| VGOS Compatibility | No | No | No | No |
| Data Input Connection | VLBA | VLBA | VSI-H | VSI-H |
| Max data-rate per input connection | 1 Gbps | 1 Gbps | 1 Gbps | 2 Gbps |
| Max aggregate input data rate | 1 Gbps | 1 Gbps | 1 Gbps | 2 Gbps |
| Max sustained record rate | 1 Gbps | 1 Gbps | 1 Gbps | 2 Gbps |
| RAM buffer size | N.A. | N.A. | N.A. | N.A. |
| Removable media? | Yes | Yes | Yes | Yes |
| Recording media packaging | Mark 5 8-disk PATA or SATA module | | | |
| Max # of media packages per system | 2 | 2 | 2 | 2 |
| # media packages for max recording data rate | 1 | 1 | 1 | 1 |
| Input data format | VLBA/Mark 4 | VLBA/Mark | VSI-H | VSI-H |

| | | | | |
|---|--|---------------------------|----------|----------|
| Record-data format | VLBA/Mark 4 data frame format | Mark 5B data frame format | | |
| Media file-system format | Proprietary (Conduant) | | | |
| Resilient to slow recording media? | Yes | Yes | Yes | Yes |
| Max aggregate playback rate | 1 Gbps | 1 Gbps | 1 Gbps | 1 Gbps |
| Playback data interface type | GigE | GigE | GigE | GigE |
| Playback data format | VLBA | VLBA | Mark 5B | Mark 5B |
| OS type | Linux | Linux | Linux | Linux |
| Software developer ('proprietary' or 'open') | Conduant (proprietary); Haystack (open) | | | |
| Command interface connection | GigE | GigE | GigE | GigE |
| Command language | VSI-S | VSI-S | VSI-S | VSI-S |
| Current status | obsolete | obsolete | obsolete | Obsolete |
| Cost (USD) | N.A. | N.A. | N.A. | N.A. |
| Contact | mark5@haystack.mit.edu | | | |

Table 3a. Legacy Mark 5 Data Recorders – not compatible with VGOS and obsolete.

VLBI Data Recorders that are not VGOS compliant, Table 3b.

LBADR. The LBADR is a data recorder designed for the Australian LBA. It uses VSI-H data interfaces and Linux files with 'scattered' data. It is packaged as a Standard PC (with optional disk array) and uses standard PC electronics plus a PCI data-capture board from Metsahovi (which is no longer available). LBADR data rates are too low to be useful for VGOS.

K5/VSSP32. The K5/VSSP32 is a highly COTS recorder. Its input connection is USB, its file system is linux, and it uses a standard COTS PC system with NIC disk interface. It does not support the standard IVS VSI-S control protocol but is controlled directly by logging into the system. K5/VSSP32 data rates are too low to be useful for VGOS.

| | <i>LBADR</i> | <i>K5/VSSP32</i> |
|---|-------------------------------|-------------------|
| VGOS Compatibility | No | No |
| Data Input Connection | 2xVSI-H | USB |
| Max data-rate per input connection | 512 Mbps | 256 Mbps |
| Max aggregate input data rate | 512 Mbps | 256 Mbps |
| Max sustained record rate | 512 Mbps | 256 Mbps |
| RAM buffer size | N.A. | N.A. |
| Removable media? | ? | Y |
| Recording media packaging | RAID array in commercial pkg. | PC removable disk |
| Max # of media packages per system | ? | Flexible |

| | | |
|---|--|---|
| # media packages for max recording data rate | 1 | 1 |
| Input data format | ? | K5-VSSP32 |
| Record-data format | ? | K5/VSSP * |
| Media file-system format | Linux files w/'scattered' data | Linux files |
| Resilient to slow recording media? | Y | Y |
| Max aggregate playback rate | ? | Depends on speed of PC ~6 Gbps |
| Playback data interface type | GigE | 1GigE or 10GigE |
| Playback data format | VSI-H or optional software defined through 1GigE | 1GigE or 10GigE |
| OS type | Linux | Linux |
| Software developer ('proprietary' or 'open') | ? | Nitsuki (prop) NICT (open) |
| Command interface connection | ? | GigE |
| Command language | ? | Direct control by logging into the system |
| Current status | PCI data-capture board no longer available | Available |
| Cost (USD) | N.A. | Price of a PC with USB |
| Contact | Chris Phillips <chris.phillips@csiro.au> | Mamoru Sekido <sekido@nict.go.jp> |

* K5/VSSP Data Format: http://www2.nict.go.jp/aeri/sts/stmg/K5/VSSP/vssp32_format.pdf

Table 3b. Other VLBI Data Recorders – not compatible with VGOS