

## An Almost Fair Comparison Between Earthworm and SeisComp3

by Marco Olivieri and John Clinton

### INTRODUCTION

Any earthquake service that maintains a seismic network operates an earthquake monitoring system (EMS), which is a combination of software capable of acquiring, processing, and archiving large volumes of real-time continuous seismic data. Crucially, the processing includes the automatic detection, location, and quantification of earthquakes. Manual review and alert dissemination are typically also core components of an EMS. There are a small number of modern and complete EMSs available to the seismological community; only a couple of these are free. A new seismic network must decide which EMS to operate, and even established networks are periodically required to consider either updating or overhauling their operational EMS. Often, a seismic network has operated a particular EMS for several years and, in some cases, decades. Although there are few options, the selection of the appropriate EMS to operate is not at all a trivial choice.

The enduringly popular Earthworm was the first true EMS. Carl Johnson from the U.S. Geological Survey (USGS), a lead developer of the original system, followed a set of design principles that “guide the design and implementation of a seismic processing system”: modularity, system independence, scalability, connectivity, and robustness (from [www.earthwormcentral.org/](http://www.earthwormcentral.org/)). More than two decades later, these could still serve as the guiding principles for the general design and development of a sustainable EMS.

In this opinion piece, we compare two of the leading cost-free options: Earthworm and SeisComp3. This opinion paper is intended to aid network seismologists who need to select or migrate to a new monitoring system for earthquake surveillance. We do not pretend to have the definitive answer to this complex decision. Indeed, selecting the right software is highly network dependent, requiring reflection on the specific network goals, existing network boundary conditions (e.g., field sensors), and available expertise. Nonetheless, we aim to offer a summary of criteria to consider and a critical evaluation of two of the best options, reflecting our joint experience. Between us, we have operated three of the largest permanent European seismic networks: the dense but small aperture Swiss Seismic Network, operated by the Swiss Seismological

Service (SED) at ETH Zurich; the MedNet Network, which covers the entire Mediterranean region; and the Italian National Network. The latter two are both operated by Istituto Nazionale di Geofisica e Vulcanologia (INGV) in Italy. In these positions, we have stretched various EMSs to fit evolving requirements, and we were eventually faced exactly with the question of how to move from a proprietary EMS to a community standard. In this last process, we gained experience not only from the available technical literature but also from many discussions with software developers and seismologists closely associated with the monitoring of regional seismicity and with experience satisfying the requirements of the relevant funding authorities, including civil protection agencies (CPAs).

In particular, this opinion paper results from some months of discussion between the two authors while setting up the roadmap for upgrading the monitoring system at the SED. The comparison cannot be completely objective, as any such effort results in a decision being made with the majority of subsequent effort focused on implementation and optimization of the preferred solution. Nevertheless, we make our best effort to be fair to both potential users and the existing user and software development communities.

Why do we compare just Earthworm and SeisComp3 and not other popular software like Antelope or Hydra or SEISAN? It is primarily because these are the two that we have experience with and also because these are the two major free, open (or almost open) source, and complete systems, each supported by strong user and active development communities. In addition, Earthworm can be viewed as the stable, robust EMS, the staple of many networks in the United States, Europe, and worldwide, whereas SeisComp3 is a young, dynamic system currently gaining popularity but as yet relatively untested. Both systems fulfill Carl Johnson’s still relevant key criteria for an EMS, although SeisComp3 is not system independent.

### KNOW WHAT YOU WANT, AND PLAN HOW TO GET THERE!

When choosing a new dishwasher, to select the appropriate model from the many on offer, we ask ourselves some simple questions. Which model is within our budget? Which brand do we trust and hence expect to operate over many years without regular servicing? Which model has all the features we need (or even just desire)? Which will fit into the slot left free after we remove the previous model? When choosing a new monitoring system, we seek answers to similar basic questions:

Which EMS will be the most cost-effective and durable in the long term? Which system can be easily installed, achieving the minimum acceptable operational setup using the available resources? Which EMS will provide the most durable system in the long term? Which offers the best feature set, including, for example, graphical user interfaces (GUIs) for manual event review? Which EMS will best match the current infrastructure, for example, field hardware or alarming software? Additional critical questions to consider include the following: What kind of seismicity do we need to monitor in terms of the geographic and magnitude distribution? Are the network density and quality sufficient to reach this goal, irrespective of the EMS used? What field hardware do we operate now and are likely to use in the future? Which of these are immovable boundaries (seismic stations, software, communication systems, alerting procedures) and which can be modified during the course of the migration? What is our commitment to our CPA or national media? This last point is important to reflect on further. Funding authorities often include both agencies involved in civil protection and science. Both are generally interested in ensuring that they have access to high-quality earthquake information, and so an EMS must be at minimum capable of providing this information. For civil protection agencies, this typically further includes a mandate to detect all events of significance and provide rapid, accurate event locations with consistent magnitude estimates, preferably with some interpretation of the significance of the event, all with a minimal number of false notifications. For scientific agencies, focus is less on rapid notification but rather on supporting research, namely, producing and maintaining a high-quality seismicity catalogue with both a low magnitude of completeness and consistently determined locations and magnitudes. Good support tools for waveform, event catalogue, and metadata dissemination are also expected.

These general questions are not new. Ottemoeller and Havskov (2011) recently posed similar questions regarding general seismic network operation. Furthermore, many of these questions often have obvious answers, especially for experienced network managers. In any case, failing to clarify the requirements of the production EMS or underestimating the complexity of the task will lead to delays and redesign or, worse still, abandonment of the project. Securing all the desired features may not be possible by simple configuration tuning or minor software development. Prior to selecting the new system, it is critical to know the human resources you have available for the installation, configuration, maintenance, or software development and understand whether, if these prove insufficient, there could be short-term funds available for having local staff or hiring external expertise. This is important because the perfect EMS does not exist, and even minor software development to meet immovable local requirements, such as implementing your network's historic Richter magnitude definition or configuring alerts to be in the format required by the local authorities, can require significant time investment for untrained staff or, alternatively, expensive contracts with relevant experts.

It is important to know the ideal acquisition—processing—archiving—dissemination scheme you wish to install and then understand what components can and cannot be realized with the EMS options under consideration. Lessons can come from other seismic networks, although different networks can attribute very different meanings to concepts like a “good location,” “completeness of detected events,” or indeed an “efficiently operating software.” On this latter point, a good EMS not only rapidly and accurately locates and quantifies the nearby  $M$  6.0 but also restarts without chaos following a software crash.

## EARTHWORM OVERVIEW

Drawing on core developments such as the Allen picker (Allen, 1982), the Earthworm project began in the early 1990s at USGS Menlo Park in the United States (Johnson *et al.*, 1995). Currently, Earthworm is a robust and well-tested software that includes exhaustive and updated documentation (see [www.earthwormcentral.org/](http://www.earthwormcentral.org/), also mirrored at <http://folkworm.ceri.memphis.edu/ew-doc/> and [www.isti.com/products/earthworm](http://www.isti.com/products/earthworm)). The core programs are written in C by seismologists. A broad community of expert users share experiences and post problems, solutions, new developments, and user suggestions via an active mailing list. A private software firm, Instrumental Software Technologies, Inc. (ISTI; [www.isti.com](http://www.isti.com)), and researchers at Center for Earthquake Research and Information (CERI; University of Memphis) jointly coordinate the development, maintain the website and mailing list, and have recently improved the tools for developers including access to a dedicated mailing list and a versioned code repository. Several research institutions and networks develop their own Earthworm code extensions that are included in the distribution or added to the section on contributed software of the website. The current general release version (v7.5) does not include interactive GUIs for catalogue cleaning, event relocation, waveform browsing, or quality control (QC) checking. This suite is an extension of the Earthworm component previously known as Automatic Earthworm. In the recent past, Earthworm releases included two versions, Automatic Earthworm and Interactive Earthworm (Earle *et al.*, 2003). The former is sufficient for automatic event detection, quantification, and alerting; the latter, which requires the commercial database software Oracle, allows interaction with the automatic location, preparation of webpages, and manual relocation of events via the external java software jiggle (<http://pasadena.wr.usgs.gov/jiggle/>). The Earthworm community expects that the currently unavailable Interactive Earthworm will be replaced by the ANSS Quake Monitoring System (AQMS; Friberg *et al.*, 2010) system, which is currently operational at many U.S. universities. AQMS also requires using an Oracle database and is not yet an open release; hence, we cannot comment further. However, extensive documentation is available at <http://vault.gps.caltech.edu/trac/cisn>. The Google group ANSS-AQMS (<http://groups.google.com/group/anss-aqms>) hosts the community mailing list.

ISTI offers a number of management GUIs that work with v7.5, such as Earthworm Front End (EWFE) and Seis-NetWatch, but these are not part of the core distribution.

Some of the Earthworm architectural information technology (IT) concepts could be considered dated by software engineers (e.g., there is no database at its core), but it continues to be operated successfully by many networks in some of the most seismically active regions in the world. The typical operator is monitoring seismicity at local or regional scales (users include, for example, the majority of regional seismic networks in the United States, INGV Italy, and Puerto Rico). The software has also been successfully extended to volcano observatories (Hawaii Volcanoes Observatory and Alaska Volcano Observatory [AVO]; Dixon *et al.*, 2005) and geomagnetic observatories, and a variant, Earlybird, is also used in many tsunami monitoring centers (Luckett *et al.*, 2008).

## SEISCOMP3 OVERVIEW

The SeisCompP project has been developed over the last decade. A major effort to consolidate existing robust software modules (e.g., SeedLink; Hanka *et al.*, 2000) into a fully functional EMS (what is now known as SeisCompP3 or third-generation SeisCompP) was completed in 2008 by the GEOFON group at GFZ Potsdam in Germany. Support for the project came from internal GFZ funds, European projects, and the German Indonesian Tsunami Early Warning System (GITEWS) project ([www.gitews.de](http://www.gitews.de)). Currently, SeisCompP3 is a complete suite of software based on a unique database and provides modules for data acquisition, data processing, archiving, and dissemination and also includes a suite of GUIs. It is a coherent software engineering effort, primarily coded in C++, with most library functionality accessible from Python scripts through a thin wrapper layer, and adheres to community standards where available, such as QuakeML (<https://quake.ethz.ch/quakeml>) and SEED ([www.iris.edu/manuals/SEEDManual\\_V2.4.pdf](http://www.iris.edu/manuals/SEEDManual_V2.4.pdf)). SeisCompP3 is rapidly gaining popularity, particularly in Europe, where components have been developed with the support of European project funding, and East Asia, where it has wide exposure as the software developed for rapid warning of large events in Indonesia. As it is relatively young, software features change relatively rapidly; documentation is not always exhaustive; and although a mailing list exists (see [www.seiscomp3.org](http://www.seiscomp3.org)), there are few experts capable of answering complex mailing list questions. A commercial software firm, GEMPA ([www.gempa.de](http://www.gempa.de)), set up by some of the core GITEWS developers, is heavily involved in software development, distribution, and maintenance. Funding for development comes from GFZ, IRIS, and the user community.

A major advantage with SeisCompP3 is that it is developed over a short period, so there is a coherent IT base. The main SeisCompP3 development in terms of the event detection and location software was driven by the GITEWS project and focused on optimizing the ability to monitor and rapidly characterize potentially tsunamigenic earthquakes at the regional and global scale (Hanka *et al.*, 2010). Currently, it is being

installed and evaluated at many regional and local networks, and hence, its ability to monitor local seismicity is rapidly being improved. It requires substantial software expertise for code bug fixing and development. In addition to GFZ, SeisCompP3 is currently operational in Indonesia and Greece and under test or evaluation in some seismically active countries (e.g., New Zealand). Some institutions have begun developing their own code and contributing modules, although currently there are limited resources for independent developers.

## DIRECT COMPARISON OF KEY FEATURES

### Start-Up Costs

The start-up costs—the time and manpower required to set up a fully functioning and operational system—are essential for a network to correctly evaluate. They are heavily dependent on the local network conditions, including the type of seismicity to be monitored, the instrumentation that must interface with the monitoring system, the number and expertise of the available network group, and the requirements imposed on the network. Seismologists are required to define the scientific requirements and optimize seismological configuration. Software engineers are needed to set up the IT framework to operate the software with the required degree of speed, robustness, and redundancy and to extend the software so that it can support the network requirements. The SED, for example, does not operate a 24-h in-office service and distributes automatic locations without review. Thus, both a robust IT setup and confidence in the automatic event identification are required. To harden the IT performance, the SED operates duplicate physical machines in separate locations running identical acquisition and processing software, contracts the university IT services to host the machines (providing backup power and communications) and to manage the archives, and has external monitoring software constantly checking for a wide variety of potential hardware and software problems, sending pager messages upon occurrence.

Both SeisCompP3 and Earthworm are relatively simple to get running. SeisCompP3 has an out-of-the-box default configuration that accesses open global stations and hence can give the impression of a plug-and-play system. This is an advantage in the short term because this configuration should provide a good location and magnitude for the next teleseismic  $M$  6.5 event, although it can also lead to complacency and a reluctance to fully comprehend the software and optimize the configuration before a final transition. It is trickier to customize the system to include monitoring of local stations. If your network maintains dataless SEED or a pole-zero file library, the entire network station information can be easily uploaded to the SeisCompP3 database. Earthworm comes with prebuilt binary bundles or it can be compiled, and local network configuration is derived from a pole-zero flat-file library that may also be generated from dataless SEED.

The true estimate of the start-up costs involves meeting the full requirements of the network. If your network is required to monitor vigorous induced sequences with  $M_L \sim 1.5$ ,

and the software has never successfully performed this task at other networks of similar station density and quality, the effort to reach this level of precision cannot be known. In these circumstances, the initial momentum gained by being successful in turning on the software can easily be lost in the winding path to get the desired results, resulting in a loss of enthusiasm and, more critically, risk of not being able to complete the project without requiring additional resources. Earthworm has the larger user community and has been used in a wider variety of monitoring contexts and hence may be more likely to provide a proven methodology if an unusual set of monitoring requirements needs to meet.

### Information Technology

Seismologists often tend to underestimate the IT costs for the design, implementation, and maintenance of a robust instance of their chosen EMS that meets local requirements. Hence, it is mandatory to have IT colleagues with an overview of the EMS software capabilities, the network's redundancy and high availability requirements, and an honest appraisal of the existing IT capabilities of the network. A key part of the IT strategy is to operate the system on an environment the network staff are comfortable operating. In this manner, Earthworm is far more flexible because it can be run on multiple operating systems (Windows, Solaris, Mac OS X, and Linux), whereas SeisComp3 runs on only various distributions of Linux. It is noted though that not all community-contributed Earthworm software operates on all platforms.

The automatic earthquake detection and quantification processing system is usually located downstream of a set of servers and services dedicated solely to data acquisition and data archiving. This automatic software can use the same infrastructure as the tools used for manual event review and dissemination. In the case of SeisComp3, because a manual review GUI exists as part of the package, using the same server for both automatic analysis and review is particularly common. Users of Earthworm, who do not run either an older version that includes Interactive Earthworm or the new AQMS (not on open release), need to design their own interface between the automatic location information and their chosen manual relocation software (this is not a trivial issue).

Both Earthworm and SeisComp3 are scalable systems that can operate on both large and small networks, capable of acquiring and processing many hundreds of high-sample-rate data streams and processing vigorous aftershock sequences. As a newer system, SeisComp3 does not include many examples of efficient operation during a robust local earthquake sequence, although it has been in real-time use during the Christchurch sequence in New Zealand (G. Clitheroe, personal comm., 2011) and tested retroactively on the L'Aquila sequence at the SED. Additional observations of behavior during robust local earthquake sequences will increase the community confidence in the software and lead to improved tuning for these critical scenarios.

For both systems, hardware should be properly dimensioned, and, for reasons ranging from redundancy, to stability,

to scalability, dividing jobs onto different machines or duplicating setups should be considered. As an example, to minimize the risk of failure on the server performing the automatic processing, the SED operates two near-identical versions of SeisComp3 sharing a common database. One machine is dedicated to automatic operations, whereas the second is used for all manual interaction, including manual relocations, moment tensor review, and catalogue management.

Neither SeisComp3 nor Earthworm offers a complete solution for security and high availability; however, on both systems, situation-specific solutions can be designed and configured by local IT experts. External system monitoring software such as Hobbit/Zymon or Nagios can be configured on either system to externally monitor and send alerts for a wide range of seismological and IT failure indicators, including interruption in the flow of real-time waveform data, the production of new picks and origins, and problems in log files.

Earthworm has defined standards for contributing code to the core software. The SeisComp3 community is currently building these standards. Many groups have contributed code to the Earthworm community; only a few are so far active within SeisComp3.

Robustness and stability are generally not an issue for both EMSs that typically run without failing as long as disk space is managed. Earthworm has the statmgr module that keeps the modules alive, restarts dead modules, and can send alerts if the module needs human intervention. In addition, the ISTI product EWFE GUI monitors the Earthworm module and messaging status. SeisComp3 lacks this feature, only having a cronjob checking module status and restarting failed modules, although no warning is issued when modules fail or indeed fail to restart. The SeisComp3 log files though can be locally parsed to send messages outside the SeisComp3 system.

The scqc module monitors waveform QC in SeisComp3, checking standard parameters like channel latency, gaps, root mean square, and offset, with results stored in the SeisComp3 database. A display GUI exists, although messages are not currently distributed. Earthworm does not come with a QC module, but the ISTI product SeisNetWatch monitors Earthworm acquisition, including the state of health channels from a wide number of dataloggers. SeisNetWatch standard distribution is free and open source. Data are stored in a dedicated database.

### Acquisition and Data Access

Both systems are capable of acquiring data from a variety of popular datalogger models, although they use different approaches. SeisComp3 acquisition is based on the SeedLink protocol. SeedLink includes a wide variety of plug-ins capable of acquiring data from the majority of common dataloggers. Data are then accessed by modules through a TCP/IP connection. The Earthworm distribution includes a variety of modules for acquiring data from the majority of popular dataloggers and datasers, including slink2ew, which supports Earthworm acquisition from SeedLink directly from the datalogger or from a SeedLink server. Earthworm uses a concept of waveform rings and messages, and continuous data are passed to



the dedicated WinstonWaveServer. This wave server can also be used as the permanent waveform archive. Earthworm processing modules, using a TCP connection to a dedicated server, connect to the wave server to access real-time or archive waveform data. In SeisComp3, SeedLink is used for real-time and near-real-time data access, and ArcLink is used to access archived data from a flat-file directory structure. In this case, the two systems share the same concept for data requests and the same gap problems in case of restarts.

Whatever EMS is used, gap handling and latency need to be understood. The network manager should know the compatibility of his or her field dataloggers or datasever and the chosen EMS acquisition software. For example, some manufacturers support SeedLink protocol at the datalogger level and thus are compatible with SeisComp3 and can also be efficiently imported to Earthworm by `slink2ew`. Other dataloggers require third-party software, running in the field either as a plug-in or on a separate PC or downstream of an acquisition server running proprietary software. In these cases, the data are accessible to Earthworm or SeisComp3 with added latency, but more critically, inevitable gaps in the real-time data flow may not be efficiently handled. These gaps typically are produced by communications problems, but they can also be due to station hardware or software issues. These gaps often occur during earthquake sequences because the communication systems are stressed and the data volume rises due to the large high-frequency amplitudes being recorded, and they lead to a reduced ability to locate and characterize the earthquakes at the most critical time. Nevertheless, if there is no sufficient local data storage capacity and good compatibility between the datalogger, the intermediate software, and the EMS, rapid and complete gap recovery can be a challenge, and in many cases, data (even from large earthquakes) are unnecessarily lost.

### Seismology for Automated Processing

Earthworm includes many of the seismological *de facto* standards for automated data processing. By standards, here, we are referring to well-tested and popular algorithms. Major advances in automated earthquake detection software were made concurrently and within the context of the Earthworm software development. For example, standard software such as the Allen picker (Allen, 1978), binder (Dietz, 2002), and Hypoinverse2000 (Klein, 2002) are part of the core Earthworm distribution. These programs have been running at many centers over decades, and the community can assume a level of algorithm robustness not generally applicable to other software. Nonetheless, modules such as the picker are complicated to optimally tune, although good documentation exists (Pechmann, 1998; Mele *et al.*, 2010). The binder is the standard tool for associating picks and attempts to solve the problem of fake associations from high-frequency teleseismic *P*-wave arrivals. Hypoinverse is also a community standard, although many other earthquake location software tools have been produced over the last 20 years.

The standard Earthworm modules for picking and location are `pick_ew`, `pkfilter`, `binder_ew`, and the “sausage” mega

module that includes `eqbuf`, `eqverify`, `eqproc`, and `hypo2000_mgr`. Alternative modules for event locations are `raypick` and `rayloc`, which are, respectively, the global picker and the relocater developed by Ray Buland, and `Glass`, which is the global earthquake associator developed by Carl Johnson. These alternative modules are adopted from the seismological core of the Hydra EMS developed and used at the National Earthquake Information Center in Colorado for its global earthquake monitoring. `Glass` can also read standard `pick_ew` picks. An additional Earthworm module provides an interface with `NonLinLoc` (Lomax *et al.*, 2000) and user-defined 1D or 3D velocity models.

Another noteworthy Earthworm event detection feature is the modules `carlStaTrig` and `carlSubTrig`, which find coincidence triggers at subsets of stations. This is particularly useful for networks covering large regions with dense subsets of stations, such as in volcano monitoring.

SeisComp3 currently supports a similar automatic detection software. The primary STA-LTA detector `scoutpick` can run alone or can be followed by `postpicking` using a variety of other pickers (recent implementations include the Baer and Kradolfer [1987] picker or the AIC picker [Leonard and Kennett, 1999]); the subsequent module `scoutloc` includes both phase association and a primary location using a `LOCSAT` (Bratt and Bache, 1988) grid search and user-configurable travel-time tables. Once a preliminary location is available, `screloc` can produce refined locations using `NonLinLoc` and user-defined 1D or 3D velocity profiles. In SeisComp3, the automatic locator assumes all the picks to be the first arriving *P* phase, whereas the Earthworm binder allows arrivals to be associated to defined *P* and *S* phases.

In SeisComp3, there are limitations in the automated processing capability for the secondary streams at stations that operate collocated sensors. In the common case of a collocated high-gain velocity sensor alongside a low-gain strong-motion accelerometer, the optimal usage for automated event characterization is not possible: The velocity sensor should be used for detection, and the strong-motion sensor should be used for magnitude estimation if the high-gain sensor saturates during a large earthquake. The manual relocation GUI `scolv` does provide this capability. Default magnitudes provided by Earthworm are  $M_D$  and  $M_L$ ; SeisComp3 provides  $M_L$  (using either horizontal or vertical components),  $M_S$ ,  $M_b$ , and  $M_w$  (derived from  $m_b$ , see Bormann and Saul, 2008).

Additional seismological information, such as the Earthworm modules that compute ground-motion parameters compatible with `ShakeMap`, or moment tensor calculations are desirable features. SeisComp3 does not include these in the latest release, although these modules are under development.

### GUIs and Database

An important advantage of SeisComp3 is its use of a single database (supporting open source options, e.g., MySQL or PostgreSQL), which uses a data model based on the `QuakeML` standard (Schorlemmer *et al.*, 2011). SeisComp3 also includes a handful of coordinated GUIs with similar layout and

consistent style, including hot keys. Although not part of the current release, Interactive Earthworm also uses a database, the commercially licensed Oracle (this is because major Earthworm developments have been made at U.S. institutions and at USGS, which have institutional access to Oracle), and the majority of existing Earthworm GUIs all require access to the Interactive Earthworm database. In contrast, SeisComp3 is a complete suite, with high-quality GUIs designed to serve both the control room in Indonesia dedicated to potential tsunamigenic large earthquakes and the local seismic network dedicated to careful review and analysis of microseismicity. Note though that the GUIs in SeisComp3 are distributed as binaries and are not open source.

A key feature of SeisComp3 is *scolv*, a GUI that allows interaction with and management of the event catalogue; access to all origins comprising each event; manual repicking of data; manual relocation of events; magnitude review using the processed waveforms; and rapid solution visualization such as station residual plots, move-out curves, and first motion plots. Other GUIs are available that show the most recent event summaries, real-time waveforms, and QC summaries.

In addition to the network monitoring software GUI SeisNetWatch, Earthworm users also can use the contributed java GUIs SWARM and pickewAnalysis. SWARM, developed at the AVO ([www.avo.alaska.edu/Software/swarm/](http://www.avo.alaska.edu/Software/swarm/)), supports the visualization and interaction with real-time data on the wave server. pickewAnalysis (authored by Ruben Luis at the University of Azores) is used to visualize and optimize the configuration of the Allen picker.

## Archiving

It is important to consider whether your EMS will also manage the archiving of continuous data. A related question is how can the manual event review software access archived waveforms. When Earthworm was first designed, storage of continuous data was extremely rare, so detected events were extracted and archived and the rest of the data deleted. Earthworm can archive event data, although because large volume storage is rapidly becoming cheaper and easier to manage, this is becoming an obsolete practice. Modern dataservers, including the SeisComp3 ArcLink software, can extract event-based data on the fly (e.g., [www.webdc.eu](http://www.webdc.eu), [www.seismicportal.eu](http://www.seismicportal.eu)) from continuous archives with only a short additional delay.

Now, continuous storage is commonplace and indeed is assumed for the SeisComp3 system. In SeisComp3, permanent archiving can be done via SeedLink using *slarchive*, and subsequent data access using GUIs or offline processing uses an automated switching between SeedLink for near-real-time data access and ArcLink for archived data access. Events can be extracted from the continuous archives on demand. Nonetheless, triggered, noncontinuous event data from strong-motion stations can be added to the continuous archive structure and can be accessed during a manual event review. Recent releases of Earthworm now support data archiving in similar file structures using *ew2mseed*. An additional option is to use the

WinstonWaveServer; its MySQL database can hold large archives depending on the available disk space.

The creation of data archives that use formats and structure compatible with the selected EMS is important beyond simply reviewing old events. Most EMSs, including SeisComp3 and Earthworm, support playback of waveforms and picks into the system (emulating a real-time data flow using archived data) so that configuration can be optimized, and network failures can be replayed with a different configuration to check whether fixes work effectively.

## Dissemination Tools: Alerts and Waveform Data

Alert dissemination is often the most network-specific part of an EMS. The exact local requirements strongly depend on the agreements between and expectations from the seismic network and the relevant local and national authorities. Such national differences occur in Italy and Switzerland. In Italy, the CPA requires a manually reviewed event notification by phone within minutes, which can only be met by operating a 24-h onsite shift. Subsequent alerting of the event to the media and the public is then through the CPA. In Switzerland, the initial alerts to the public, media, and authorities use automatic solutions and are delivered (in four different languages) via e-mail, SMS, FTP, and the web. The on-call duty seismologist only reviews the automatic location from home if the event occurs outside of office hours. Each agency thus has a different perspective when creating not only their alerting mechanisms but also their event review procedures and defining their tolerance to false automatic alarms.

Because of these specific local requirements, networks often develop and maintain their own alerting software. Nonetheless, SeisComp3 includes modules allowing rapid immediate review of automatic events (*scolv*) and sends messages (e-mails, SMS, voice announcements using *scvoice/scalert*) after an event that reaches a configurable alarming criterion, such as a new earthquake, an updated location, or high amplitudes at a particular station. Earthworm does not have such a core tool, although several users have developed their own codes for review and alert, which are included in the contributed code distribution.

A major advantage of SeisComp3 is its embedded distribution software ArcLink; SeisComp3 thus supports not only real-time data sharing between networks via SeedLink but also public, scientific, and cross-network access to continuous waveform archives via ArcLink (see, e.g., the web interface [arclink.ethz.ch](http://arclink.ethz.ch) and the command-line tool `arclink_fetch`). ArcLink is also the backbone distribution software behind the European Integrated Data Archive (e.g., [www.webdc.eu](http://www.webdc.eu) or [eida.ethz.ch](http://eida.ethz.ch)) and is used at the European Seismic Portal ([www.seismicportal.eu](http://www.seismicportal.eu)). For the Earthworm community, *ew2mseed* can be used in association with the WinstonWaveServer to share archive data.

## Documentation and Logging

Currently, there is no competition between Earthworm and SeisComp3 on this topic. Earthworm's documentation is complete and exhaustive; the SeisComp3 documentation is

still under development and not considered the highest priority because software development continues rapidly. Critical modules, such as the associator and the locator `scoutoloc`, are not exhaustively documented, making it difficult to track and interpret unexpected behavior. Both EMSs have online documentation. SeisComp3 uses a WIKI, allowing community-contributed updates ([www.seiscomp3.org/](http://www.seiscomp3.org/)). Furthermore, the Earthworm code is well commented; in contrast, parts of the SeisComp3 code are uncommented. Both the systems use configurable logging levels, very useful when debugging, and the verbosity helps in tracking problems.

Improved documentation and usability come with having a large, enthusiastic user community, and SeisComp3 is only at the early stages in this regard.

### License

Following USGS policy, the Earthworm core is distributed without any license. Some contributed codes (and not all) are licensed under GPL (Generic Public License). SeisComp3 has three different licenses for the different parts of the code: GNU, SeisComp public, and binary. The public license covers the seismological and infrastructural core; the binary license covers the majority of the GUIs, which are not open source; and the GNU license contains SeedLink and other libraries. Using the GUIs under the binary license requires a signed agreement with GFZ. The SeisComp3 community would benefit from a simpler licensing scheme providing more freedom to the users and developers.

### SUMMARY AND FUTURE OUTLOOK

Earthworm is a stable and relatively bug-free system in development since the early 1990s. SeisComp3 is a recently developed complex suite with major advantages including coordinated software development of all core features, integrated GUIs, and a single underlying database. Earthworm is commonly operated under Solaris, Windows, Mac OS, and Linux, whereas SeisComp3 runs only on various Linux flavors. Although both EMSs grow by community contributions, each is underpinned by commercial software companies that maintain, develop, and distribute new releases. There is thus a non-negligible possibility that if funding is tight, the long-term development and support could be compromised. Having large numbers of network seismologists and software developers working on each EMS helps mitigate this risk; a strong user and developer community improves the EMS by sharing experiences, ideas, and developments; contributing updates and bug fixing; and when necessary, providing financial support to the commercial software company for tailored installations, maintenance, and software development.

Earthworm distribution follows primarily USGS requirements, and maintainers do not always validate and include new modules in the core release; these are usually distributed in a separate contributed section. In general, the strong development community is a very positive feature, but the open contributions can limit the coherent development of the

Earthworm package. Community contributions to SeisComp3 are relatively few, and as of yet, a well-trodden path for community development has yet to emerge. Some developments have been integrated into the core, and others are distributed on the WIKI as Python scripts.

Many of the features, and indeed limitations, described in this manuscript will be out-of-date within months of publication because enhancements are introduced in both EMSs by the active development programs. Both systems will continue to perform well in terms of acquiring data, detecting local seismicity, and ensuring data archiving and storage and are likely to be among the handful of popular, reliable EMSs available to the community for years to come. ☒

### ACKNOWLEDGMENTS

Although the opinions in this review are our own, we have over the years benefited from fruitful discussions with Philip Kästli, Franco Mele, Salvatore Mazza, Doug Neuhauser, Reinoud Sleeman, Joachim Saul, Paul Friberg, Jan Becker, Geoff Clitheroe, Roman Racine, Stefan Heimers, Carlo Cauzzi, Tobias Diehl, and Alberto Michellini. We thank in particular Joachim Saul and Paul Friberg for their technical criticism and for providing the viewpoint from the SeisComp3 and Earthworm communities, respectively. Nicos Melis and Assistant Editor John Louie reviewed the manuscript, polishing the text and providing comments that substantially improved the manuscript.

### REFERENCES

- Allen, R. V. (1978). Automatic earthquake recognition and timing from single traces, *Bull. Seismol. Soc. Am.* **68**, 1521–1532.
- Allen, R. V. (1982). Automatic phase pickers: their present use and future prospects, *Bull. Seismol. Soc. Am.* **72**, S225–S242.
- Baer, M., and U. Kradolfer (1987). An automatic phase picker for local and teleseismic events, *Bull. Seismol. Soc. Am.* **77**, 1437–1445.
- Bormann, P., and J. Saul (2008). The new IASPEI standard broadband magnitude  $m_b$ , *Seismol. Res. Lett.* **79**, no. 5, 698–705.
- Bratt, S., and T. C. Bache (1988). Locating events with a sparse network of regional arrays, *Bull. Seismol. Soc. Am.* **78**, 780–798.
- Dietz, L. (2002). Notes on configuring binder\_ew: Earthworm's phase associator, [www.isti2.com/ew/ovr/binder\\_setup.html](http://www.isti2.com/ew/ovr/binder_setup.html) (last accessed May 2012).
- Dixon, J. P., J. A. Power, and S. D. Stihler (2005). A comparison of seismic event detection with IASPEI and Earthworm acquisition systems at Alaskan volcanoes, *Seismol. Res. Lett.* **76**, no. 2, 168–176.
- Earle, P., A. Bittenbinder, B. Bogaert, and C. Johnson (2003). Turn to the worm: seismic network operation using the USGS Earthworm system, in *Observations and Research Facilities for European Seismology*, *ORFEUS Newsl.* **5**, no. 1, <http://www.orfeus-eu.org/Organization/Newsletter/vol5no1/earthworm.html>.
- Friberg, P., S. Lisowski, I. Dricker, and S. Hellman (2010). Earthworm in the 21st century, *Geophys. Res. Abstr.* **12**, no. 12, 654.
- Hanka, W., A. Heinloo, and K. H. Jaekel (2000). Networked seismographs: GEOFON real-time data distribution, *ORFEUS Newsl.* **2**, no. 3, <http://www.orfeus-eu.org/Organization/Newsletter/vol2no3/geofon.html>.
- Hanka, W., J. Saul, B. Weber, J. Becker, P. Harjadi, Fauzi, and GITEWS Seismology Group (2010). Real-time earthquake monitoring for tsunami warning in the Indian Ocean and beyond, *Nat. Hazards Earth Syst. Sci.* **10**, 2611–2622.

- Johnson, C. E., A. Bittenbinder, B. Bogaert, L. Dietz, and W. Kohler (1995). Earthworm: a flexible approach to seismic network processing, *IRIS Newsl.* **14**, no. 2, 1–4.
- Klein, F. W. (2002). User's guide to Hypoinverse-2000, a FORTRAN program to solve for earthquake locations and magnitude, *U.S. Geol. Surv. Open-File Rep. 02-171*, 121.
- Leonard, M., and B. L. N. Kennett (1999). Multi-component autoregressive techniques for the analysis of seismograms, *Phys. Earth Planet. Inter.* **113**, 247–263.
- Lomax, A., J. Virieux, P. Volant, and C. Berge (2000). Probabilistic earthquake location in 3D and layered models: introduction of a Metropolis–Gibbs method and comparison with linear locations, in *Advances in Seismic Event Location*, C. H. Thurber and N. Rabinowitz (Editors), Kluwer, Amsterdam, 101–134.
- Lockett, R., L. Ottemoeller, and P. Whitmore (2008). A tsunami warning system for the northeast Atlantic, *ORFEUS Newsl.* **8**, no. 1, [http://www.orfeus-eu.org/Organization/Newsletter/vol8no1/tsunami\\_warning\\_system/tsunami\\_warning\\_system.htm](http://www.orfeus-eu.org/Organization/Newsletter/vol8no1/tsunami_warning_system/tsunami_warning_system.htm).
- Mele, F., A. Bono, V. Lauciani, A. Mandiello, C. Marocci, S. Pintore, M. Quintiliani, L. Scognamiglio, and S. Mazza (2010). Tuning an Earthworm phase picker: some considerations on the pick\_ew parameters, *Rapp. Tec. INGV* **164**, <http://istituto.ingv.it/l-ingv/produzione-scientifica/rapporti-tecnici-ingv/>.
- Ottemoeller, L., and J. Havskov (2011). Challenges when establishing a seismic network, *Seismol. Res. Lett.* **82**, 373–374.
- Pechmann, J. C. (1998). Suggested picker parameters changes, [www.isti2.com/ew/ovr/picker\\_tune.html](http://www.isti2.com/ew/ovr/picker_tune.html) (last accessed May 2012).
- Schorlemmer, D., F. Euchner, P. Kästli, and J. Saul (2011). QuakeML: status of the XML-based seismological data exchange format, *Ann. Geophys.* **54**, no. 1, 59–65.

*Marco Olivieri*  
**Istituto Nazionale di Geofisica e Vulcanologia (INGV)**  
 Bologna  
 via DonatoCreti 12  
 40128 Bologna, Italy  
 marco.olivieri@bo.ingv.it

*John Clinton*  
 Swiss Seismological Service  
 ETH Zürich  
 8092, Switzerland